




**FORTINET**®



# FortiOS™ Handbook - Troubleshooting

VERSION 5.6.4



## **FORTINET DOCUMENT LIBRARY**

<http://docs.fortinet.com>

## **FORTINET VIDEO GUIDE**

<http://video.fortinet.com>

## **FORTINET BLOG**

<https://blog.fortinet.com>

## **CUSTOMER SERVICE & SUPPORT**

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

## **FORTIGATE COOKBOOK**

<http://cookbook.fortinet.com>

## **FORTINET TRAINING SERVICES**

<http://www.fortinet.com/training>

## **FORTIGUARD CENTER**

<http://www.fortiguard.com>

## **FORTICAST**

<http://forticast.fortinet.com>

## **END USER LICENSE AGREEMENT**

<http://www.fortinet.com/doc/legal/EULA.pdf>

## **FORTINET PRIVACY POLICY**

<https://www.fortinet.com/corporate/about-us/privacy.html>

## **FEEDBACK**

Email: [techdocs@fortinet.com](mailto:techdocs@fortinet.com)



April 26, 2018

FortiOS™ Handbook - Troubleshooting

# TABLE OF CONTENTS

<b>Change Log</b>	<b>6</b>
<b>Introduction</b>	<b>7</b>
Before you begin	7
How this guide is organized	7
<b>Troubleshooting methodologies</b>	<b>8</b>
Ensure you have administrator-level access to required equipment	8
Establish a baseline	8
Define the problem	9
Create a troubleshooting plan	10
Providing supporting elements	10
Obtain any required equipment	11
Consult Fortinet resources	11
<b>Troubleshooting tools</b>	<b>12</b>
FortiOS diagnostics	12
Date and time	12
Resource usage	13
Proxy operation	14
Hardware NIC	14
Traffic trace	17
Session table	18
Firewall session setup rate	20
Finding object dependencies	21
Flow trace	22
Packet sniffing and packet capture	25
NPU-based interfaces	30
Debug command	30
The execute tac report command	32
Other commands	32
FortiOS ports	34
FortiAnalyzer and FortiManager ports	35
FortiGuard troubleshooting	35
Troubleshooting process for FortiGuard updates	36

FortiGuard server settings.....	37
<b>Troubleshooting tips.....</b>	<b>39</b>
How to check hardware connections.....	40
How to check FortiOS network settings.....	41
Interface settings.....	41
DNS settings.....	42
DHCP server settings.....	42
How to check CPU and memory resources.....	42
How to troubleshoot high memory usage.....	43
How to troubleshoot high CPU usage.....	44
How to check modem status.....	47
How to run ping and traceroute.....	48
Ping.....	48
Traceroute.....	50
How to check the logs.....	52
How to verify the contents of the routing table (in NAT mode).....	52
How to verify the correct route is being used.....	53
How to verify the correct firewall policy is being used.....	54
How to check the bridging information in transparent mode.....	54
How to check the bridging information.....	54
How to display forwarding domain information.....	54
How to check the number of sessions that UTM proxy uses.....	56
Conserve or failopen mode.....	56
How to examine the firewall session list.....	57
Check source NAT information.....	57
How to check wireless information.....	58
Troubleshooting station connection issue.....	58
Enable diagnostics for a particular station.....	58
How to verify connectivity to FortiGuard.....	58
How to perform a sniffer trace (CLI and packet capture).....	59
How do you sniff packets.....	59
How do you use packet capture.....	60
How to debug the packet flow.....	62
<b>Troubleshooting resources.....</b>	<b>63</b>
Technical documentation.....	63
Fortinet video library.....	63
Release notes.....	63
Knowledge base.....	63
Fortinet technical discussion forums.....	63
Fortinet training services online campus.....	64

Fortinet customer support .....64

## Change Log

Date	Change Description
April 20, 2018	Updated 5.6 content published.

# Introduction

Welcome and thank you for selecting Fortinet products to protect your organization's network.

This guide is intended for administrators who need guidance on different network needs and information on basic and advanced troubleshooting.

This chapter contains the following topics:

- [Before you begin](#)
- [How this guide is organized](#)

## Before you begin

Before you begin using this guide, verify the following:

- You have administrative access to the FortiGate GUI and/or CLI.
- The FortiGate is integrated into your network.
- The operation mode is configured.
- The system time, DNS settings, administrator password, and network interfaces are configured.
- Firmware, FortiGuard AntiVirus, FortiGuard Application Control, and FortiGuard IPS are up to date.

To complete the tasks in this guide, you must have super\_admin administrative access (unless otherwise specified). Users with other types of administrative access may not be able to complete the troubleshooting tasks.

## How this guide is organized

This handbook chapter presents troubleshooting and problem solving issues that may help you with your FortiGate and contains the following sections:

- ["Troubleshooting methodologies" on page 8](#) walks you through best practice concepts of FortiOS troubleshooting.
- ["Troubleshooting tools" on page 12](#) describes some of the basic commands and parts of FortiOS that can help you with troubleshooting.
- ["Troubleshooting tips" on page 39](#) presents most of the common issues and how to address them.
- ["Troubleshooting resources" on page 63](#) identifies Fortinet resources for troubleshooting.

# Troubleshooting methodologies

This section explains how to prepare for troubleshooting, create a troubleshooting plan, and where to find additional resources.

The following topics are covered:

- [Ensure you have administrator-level access to required equipment](#)
- [Establish a baseline](#)
- [Define the problem](#)
- [Create a troubleshooting plan](#)
- [Obtain any required equipment](#)
- [Consult Fortinet resources](#)

## Ensure you have administrator-level access to required equipment

Before troubleshooting your FortiGate, you need administrator access to the equipment. If you're a client on a FortiGate that has virtual domains (VDOMs) enabled, you can often troubleshoot within your own VDOM. However, you should inform the super admin for the FortiGate that you'll be performing troubleshooting tasks.

Also, you may need access to other networking equipment, such as switches, routers, and servers, to carry out tests. If you don't have access to this equipment, contact your network administrator for assistance.

## Establish a baseline

A FortiGate operates at all layers of the OSI model. For this reason, troubleshooting problems can become complex. Establishing baseline parameters for your system before a problem occurs helps to reduce the complexity when you need to troubleshoot.

Many of the guiding questions in the following sections compare the current problem to normal operation on your FortiGate. A best practice is to establish and record the normal operating status. Regular operation data shows trends, and allows you to see when changes occur and when there may be a problem. You can gather this data by using logs and SNMP tools to monitor the system performance or by regularly running information gathering commands and saving the output.



Back up your FortiOS configuration on a regular basis. This is a good practice and not only for troubleshooting. You can restore the backed up configuration as needed and save the time and effort of recreating it from the factory default settings.

---

You can use the following CLI commands to obtain normal operating data for a FortiGate:

```
get system status
```

Displays firmware versions and FortiGuard engine versions, and other system information



<code>get system performance status</code>	Displays CPU and memory states, average network usage, average sessions and session setup rate, virus caught, IPS attacks blocked, and uptime
<code>get hardware memory</code>	Displays information about memory
<code>get system session status</code>	Displays total number of sessions
<code>get router info routing-table all</code>	Displays all the routes in the routing table, including their type, source, and other useful data
<code>get ips session</code>	Displays memory used and maximum amount available to IPS as well and counts
<code>get webfilter ftgd-statistics</code>	Displays a list of FortiGuard related counts of status, errors, and other data
<code>diagnose system session list</code>	Displays the list of current detailed sessions
<code>show system dns</code>	Displays the configured DNS servers
<code>diagnose system ntp status</code>	Displays information about NTP servers

These commands are just a sample. You can run any commands for information gathering that apply to your system. For example, if you have active VPN connections, use the `get vpn *` series of commands to get more information about them.

To see an extensive snapshot of your system, you can use the `execute tac report` command. This command runs many diagnostic commands for specific configurations. Regardless of the features deployed on your FortiGate, this command records the current state of each feature. Then, if you need to perform troubleshooting later, you can run the same command again and compare the differences to quickly identify any suspicious output.

## Define the problem

The following questions help you define the problem. Be as specific as possible with your answers. Once you define the problem, you can search for a solution and then create a plan for how to solve it.

- **what's the problem?**

The problem being observed isn't necessarily the actual problem. You should determine where the problem lies before starting to troubleshoot the FortiGate.

- **Was the device working before?**

If the device never worked, it might be defective. For more information, see [Troubleshooting your FortiGate Installation](#) in the [Getting Started](#) chapter.

- **Can you reproduce the problem ?**

If the problem is intermittent, it may be dependent on system load. Note that it may be difficult to troubleshoot an

intermittent problem because it's difficult to reproduce.

- **What has changed?**

Don't assume that nothing has changed in the network. Use the FortiGate event log to identify any possible configuration changes. Changes can be in the operating environment. For example, there might be a gradual increase in load as more sites are forwarded through the firewall.

If something has changed, roll back the change and assess the impact.

- **What's the scope of the problem?**

After you isolate the problem, determine what applications, users, devices, and operating systems are affected.

- What's not working? Be specific.
- Is there more than one thing that isn't working?
- Is it partly working? If so, what parts are working?
- Is it a connectivity issue for the entire device, or is there an application that isn't reaching the Internet?
- Where did the problem occur?
- When did the problem occur and to which users or groups of users?
- What components are involved?
- What applications are affected?
- Can you use a packet sniffer to trace the problem?
- Can you use system debugging or look in the session table to trace the problem?
- Do any of the log files indicate a failure has occurred?

The answers to these questions help you narrow down the problem and identify what you should check during your troubleshooting. The more things you can eliminate, the fewer things you need to check during troubleshooting. For this reason, be as specific and accurate as you can when you gather information.

## Create a troubleshooting plan

Once you define the problem and gather facts, you can create a troubleshooting plan to solve the problem.

You should list all possible causes of the problem that you can think of and how you can test for each possible cause. The plan acts as a checklist so that you know what you've tried and what's left to check. This is also important to have if more than one person is performing troubleshooting tasks. Be ready to add to your plan, as needed.

A troubleshooting plan is also useful when you contact Fortinet Support. It provides you with a record of information about the problem and the tasks that you tried to fix it.

## Providing supporting elements

If you contact [Fortinet Support](#) to speak to the Technology Assistance Center (TAC), be prepared to provide the following information:

- Firmware build version (use the `get system status` command)
- Network topology diagram

- Recent configuration file
- Recent debug log (optional)
- Summary of troubleshooting steps that you've already taken and the results.



Don't provide the output from the `execute tac` report unless TAC requests it. The output from this command is very large and isn't required in many cases.

---

## Obtain any required equipment

To test your solution, you may require additional networking equipment, computers, or other equipment.

Network administrators usually have additional networking equipment available to loan you, or a lab where you can bring the FortiGate unit to test.

If you don't have access to equipment, check for shareware applications that can perform the same tasks. Often, there are software solutions that you can use when hardware is too expensive.

## Consult Fortinet resources

After you define your problem, create a plan to find a solution, and carry out that plan. If you can't resolve the problem, see ["Troubleshooting resources" on page 63](#).

# Troubleshooting tools

FortiOS provides a number of tools that help with troubleshooting both hardware and software issues. These tools include diagnostics and ports; ports are used when you need to understand the traffic coming in and going out on a specific port, for example, UDP 53, which is used by the FortiGate unit for DNS lookup and RBL lookup.

This section also contains information about troubleshooting FortiGuard issues.

- ["FortiOS diagnostics" on page 12](#)
- ["FortiOS ports" on page 34](#)
- ["FortiAnalyzer and FortiManager ports" on page 35](#)
- ["FortiGuard troubleshooting" on page 35](#)

## FortiOS diagnostics

FortiOS has a collection of diagnostic commands that you can use to troubleshoot and monitor the performance of your network. The `get` and `diagnose` CLI commands are the two main groups of diagnostic commands. Both commands display information about system resources, connections, and settings that allow you to locate problems or monitor system performance.

This topic includes diagnostics commands to help with:

- [Date and time](#)
- [Resource usage](#)
- [Proxy operation](#)
- [Hardware NIC](#)
- [Traffic trace](#)
- [Session table](#)
- [Firewall session setup rate](#)
- [Finding object dependencies](#)
- [Flow trace](#)
- [Packet sniffing and packet capture](#)
- [NPU-based interfaces](#)
- [Debug command](#)
- [The execute tac report command](#)
- [Other commands](#)

### Date and time

The system date and time are important for FortiGuard services, logging events, and sending alerts. The wrong time makes the log entries confusing and difficult to use.

Use Network Time Protocol (NTP) to set the date and time, if possible. This is an automatic method that doesn't require manual intervention. However, you must ensure that the port is allowed through the firewalls on your network. FortiToken synchronization requires NTP in many situations.

### How to set the date and time - GUI

1. Go to the **System Information** widget on the **Dashboard**. The date and time are displayed next to **System Time**.
2. To adjust the date and time settings, go to **System > Settings**.
3. In **System Time**, you can set the time zone, date and time, and select **NTP usage**.

### How to check the date and time - CLI

You can check the date and time using the CLI commands `execute date` and `execute time`.

```
config system global
    set timezone <integer>
end
config system ntp
    set type custom
    config ntpserver
        edit 1
            set server "ntp1.fortinet.net"
        next
        edit 2
            set server "ntp2.fortinet.net"
        next
    end
    set ntpsync enable
    set syncinterval 60
end
```

Use the `set timezone ?` command to display a list of timezones and the integers that represent them.

## Resource usage

Each program that runs on a computer has one or more processes associated with it. For example, if you open a Telnet program, it has an associated Telnet process. The same is true in FortiOS. All processes share system resources in FortiOS, including memory and CPU.

Use the `get system performance status` command to show the FortiOS performance status.

#### Sample output:

```
FGT# get system performance status
CPU states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU0 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU1 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU2 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU3 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 4050332k total, 527148k used (13%), 3381312k free (83%), 141872k freeable (3%)
Average network usage: 41 / 28 kbps in 1 minute, 54 / 44 kbps in 10 minutes, 42 / 34
kbps in 30 minutes
Average sessions: 33 sessions in 1 minute, 48 sessions in 10 minutes, 38 sessions in 30
minutes
Average session setup rate: 0 sessions per second in last 1 minute, 0 sessions per
second in last 10 minutes, 0 sessions per second in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 0 days, 22 hours, 59 minutes
```

Monitor the CPU and memory usage of internal processes, using the following command:

```
get system performance top <delay> <max_lines>
```

The data that the command lists includes the name of the daemon, the process ID, whether the process is sleeping or running, the CPU use percentage, and the memory use percentage.

**Sample output:**

```
get system performance top 10 100
Run Time: 0 days, 23 hours and 4 minutes
OU, ON, OS, 100I, OWA, OHI, OSI, OST; 3955T, 3298F
  httpsd 212 S 0.4 0.6
  forticron 169 S 0.4 0.4
  newcli 4054 R 0.4 0.2
  reportd 174 S 0.0 1.4
  pyfcgid 325 S 0.0 0.8
  cmdbsvr 141 S 0.0 0.7
  miglogd 160 S 0.0 0.6
  httpsd 211 S 0.0 0.6
  src-vis 180 S 0.0 0.6
  pyfcgid 327 S 0.0 0.6
  pyfcgid 328 S 0.0 0.6
  pyfcgid 329 S 0.0 0.6
  httpsd 162 S 0.0 0.5
  cw_acd 189 S 0.0 0.5
  httpsd 3998 S 0.0 0.5
  httpsd 4050 S 0.0 0.5
  updated 176 S 0.0 0.5
  httpsd 4052 S 0.0 0.4
  miglogd 203 S 0.0 0.4
  miglogd 204 S 0.0 0.4
```

## Proxy operation

Monitor proxy operations, using the following command:

```
diagnose test application <application> <option>
```

To display a list of available <application> values, enter:

```
diagnose test application ?
```

The <option> value depends on the application value that you use in the command. To display a list of available <option> values, enter:

```
diagnose test application <application> ?
```

For example, if the application is http, the CLI command that displays the <option> values is:

```
diagnose test application http ?
```

## Hardware NIC

Monitor hardware network operations, using the following command:

```
diagnose hardware deviceinfo nic <interface>
```

The information that this command shows is important because errors at the interface indicate data link or physical layer issues which may impact the performance of the FortiGate.

The following example shows a sample output when you set <interface> to lan:

```
System_Device_Name lan
Current_HWaddr 00:09:0f:68:35:60
Permanent_HWaddr 00:09:0f:68:35:60
State up
Link up
Speed 100
Duplex full
[.....]
Rx_Packets=5685708
Tx_Packets=4107073
Rx_Bytes=617908014
Tx_Bytes=1269751248
Rx_Errors=0
Tx_Errors=0
Rx_Dropped=0
Tx_Dropped=0
[....]
```

The `diagnose hardware deviceinfo nic` command displays a list of error names and values that are related to hardware. The following table describes possible hardware errors:

Field	Description
Rx_Errors = rx error count	Bad frame was marked as error by PHY
Rx_CRC_Errors + Rx_Length_Errors - Rx_Align_Errors	This error is only valid in 10/100M mode
Rx_Dropped or Rx_No_Buffer_Count	Running out of buffer space
Rx_Missed_Errors	Equals Rx_FIFO_Errors + CEXTERR (Carrier Extension Error Count); only valid in 1000M mode, which is marked by PHY
Tx_Errors = Tx_Aborted_ Errors	ECOL (Excessive Collisions Count); only valid in half-duplex mode

Field	Description
	Late Collisions (LATECOL) Count
Tx_Window_Errors	Late collisions are collisions that occur after 64-byte time into the transmission of the packet while working in 10 to 100 Mb/s data rate and 512-byte time into the transmission of the packet while working in the 1,000 Mb/s data rate. This register only increments if transmits are enabled and the device is in half-duplex mode.
Rx_Dropped	See Rx_Errors
Tx_Dropped	Not defined
Collisions	Total number of collisions experienced by the transmitter; valid in half-duplex mode
Rx_Length_Errors	Transmission length error
Rx_Over_Errors	Not defined
Rx_CRC_Errors	Frame CRC error
Rx_Frame_Errors	Same as Rx_Align_Errors  This error is only valid in 10/100M mode.
Rx_FIFO_Errors	Same as Rx_Missed_Errors - a missed packet count
Tx_Aborted_Errors	See Tx_Errors
Tx_Carrier_Errors	The PHY should assert the internal carrier sense signal during every transmission. Failure to do so may indicate that the link has failed or the PHY has an incorrect link configuration. This register only increments if transmits are enabled. This register isn't valid in internal SerDes 1 mode (TBI mode for the 82544GC/EI) and is valid only when the Ethernet controller is operating at full duplex.
Tx_FIFO_Errors	Not defined
Tx_Heartbeat_Errors	Not defined
Tx_Window_Errors	See <a href="#">LATECOL</a>
Tx_Single_Collision_Frames	Counts the number of times that a successfully transmitted packet encountered a single collision  The value increments only if transmits are enabled and the Ethernet controller is in half-duplex mode.



Field	Description
Tx_Multiple_Collision_Frames	A Multiple Collision Count which indicates the number of times that a transmit encountered more than one collision, but less than 16. The value increments only if transmits are enabled and the Ethernet controller is in half-duplex mode.
Tx_Deferred	Counts defer events  A deferred event occurs when the transmitter can't immediately send a packet due to the medium being busy because another device is transmitting, the IPG timer hasn't expired, half-duplex deferral events are occurring, XOFF frames are being received, or the link isn't up. This register only increments if transmits are enabled. This counter doesn't increment for streaming transmits that are deferred due to TX IPG.
Rx_Frame_Too_Longs	The Rx frame is oversized
Rx_Frame_Too_Shots	The Rx frame is too short
Rx_Align_Errors	This error is only valid in 10/100M mode
Symbol Error Count	Counts the number of symbol errors between reads - SYMERRS.  The count increases for every bad symbol that's received, whether or not a packet is currently being received and whether or not the link is up. This register increments only in internal SerDes mode.

## Traffic trace

Traffic tracing allows you to follow a specific packet stream. This is useful to confirm that packets are taking the route you expect them to take on your network.

View the characteristics of a traffic session through specific security policies using:

```
diagnose system session
```

Trace per-packet operations for flow tracing using:

```
diagnose debug flow
```

Trace per-Ethernet frame using:

```
diagnose sniffer packet
```

Trace a route from a FortiGate to a destination IP address

```
# execute traceroute www.fortinet.com
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
1 172.20.120.2 0.637 ms 0.653 ms 0.279 ms
2 209.87.254.221 <static-209-87-254-221.storm.ca> 2.448 ms 2.519 ms 2.458 ms
3 209.87.239.129 <core-2-g0-2.storm.ca> 2.917 ms 2.828 ms 9.324 ms
4 209.87.239.199 <core-3-bdi1739.storm.ca> 13.248 ms 12.401 ms 13.009 ms
5 216.66.41.113 <v502.core1.tor1.he.net> 17.181 ms 12.422 ms 12.268 ms
```

```
6 184.105.80.9 <100ge1-2.core1.nyc4.he.net> 21.355 ms 21.518 ms 21.597 ms
7 198.32.118.41 <ny-paix-gni.twgate.net> 83.297 ms 84.416 ms 83.782 ms
8 203.160.228.217 <217-228-160-203.TWGATE-IP.twgate.net> 82.579 ms 82.187 ms 82.066 ms
9 203.160.228.229 <229-228-160-203.TWGATE-IP.twgate.net> 82.055 ms 82.455 ms 81.808 ms
10 203.78.181.2 82.262 ms 81.572 ms 82.015 ms
11 203.78.186.70 83.283 ms 83.243 ms 83.293 ms
12 66.171.127.177 84.030 ms 84.229 ms 83.550 ms
13 66.171.121.34 <www.fortinet.com> 84.023 ms 83.903 ms 84.032 ms
14 66.171.121.34 <www.fortinet.com> 83.874 ms 84.084 ms 83.810 ms
```

## Session table

A session is a communication channel between two devices or applications across the network. Sessions allow FortiOS to inspect and act on a sequential group of packets in a session all together instead of inspecting each packet individually. Each of these sessions has an entry in the session table that includes important information about the session.

### Use as a tool

Session tables are useful troubleshooting tools because they allow you to verify open connections. For example, if you have a web browser open to browse the Fortinet website, you would expect a session entry from your computer, on port 80, to the IP address for the Fortinet website. Another troubleshooting method is if there are too many sessions for FortiOS to process, you can examine the session table for evidence why this is happening.

You can view the FortiGate session table from using the FortiGate GUI or the CLI. The most useful troubleshooting data comes from the CLI. The session table in the GUI also provides useful summary information, particularly the current policy number that the session is using.

### GUI session information

You can view session information by going to the **FortiView** page. Read more about [FortiView consoles](#) in the Handbook's FortiView chapter.

## How to find which security policy a specific connection is using

Every program and device on your network must have a communication channel, or session, open to pass information. The FortiGate manages these sessions with features such as traffic shaping, antivirus scanning, and blocking known bad web sites. Each session has an entry in the session table.

You may want to find information for a specific session for troubleshooting. For example, if a secure web browser session isn't working properly, you can check the session table to ensure the session is still active and going to the proper address. The session table can also tell you the security policy number it matches, so you can check what's happening in that policy.

### 1. Know your connection information.

You need to be able to identify the session you want. For this, you need the source IP address (usually your computer), the destination IP address (if you have it) and the port number which is determined by the program that you're using. Some common ports are:

- Port 80 (HTTP for web browsing),
- Port 22 (SSH used for secure login and file transfers)
- Port 23 (Telnet for a text connection)
- Port 443 (HTTPS for secure web browsing)

## 2. Find your session and policy ID.

Go to **FortiView > All Sessions**. Find your session by finding your source IP address, destination IP address (if you have it), and port number. The policy ID is listed after the destination information. If the list of sessions is very long, you can filter the list to make it easier to find your session.

## 3. When there are many sessions, use a filter to help you find your session.

If there are multiple pages of sessions, it's difficult to find a single session. You can use a filter to block out sessions that you don't want. Click the search icon on the column heading to select the filter. Select **Source IP** and enter your source IP address. Now, only sessions that originate from your IP address are displayed in the session table. If the list is still too long, you can do the same for the **Source port**. That makes it easy to find your session and the security policy ID. When you're finished, clear the filters.

## CLI session information

The session table output that the `diagnose system session list` command generates is very large. You can use filters to display only the session data that you're interested in.

An entry is placed in the session table for each traffic session passing through a security policy. The following command lists the information for a session in the table:

```
diagnose system session list
```

The filter option displays specific information, for example:

```
diagnose system session filter <option>
```

The values for <option> include the following:

<code>clear</code>	Clear session filter
<code>dintf</code>	Destination interface
<code>dport</code>	Destination port
<code>dst</code>	Destination IP address
<code>duration</code>	Duration of the session
<code>expire</code>	Expire
<code>negate</code>	Inverse filter
<code>nport</code>	NAT'd source port
<code>nsrc</code>	NAT'd source IP address
<code>policy</code>	Policy ID
<code>proto</code>	Protocol number

<code>proto-state</code>	Protocol state
<code>session-state1</code>	Session state1
<code>session-state2</code>	Session state2
<code>sintf</code>	Source interface
<code>sport</code>	Source port
<code>src</code>	Source IP address
<code>vd</code>	Index of virtual domain, -1 matches all

Even though UDP is a sessionless protocol, the FortiGate keeps track of the following states:

- When UDP reply doesn't have a value of 0
- When UDP reply has a value of 1

The following table displays firewall session states from the session table:

State	Description
<code>log</code>	Session is being logged
<code>local</code>	Session is originated from or destined for local stack
<code>ext</code>	Session is created by a firewall session helper
<code>may_dirty</code>	Session is created by a policy For example, the session for <code>ftp control channel</code> will have this state but <code>ftp data channel</code> won't. This is also seen when NAT is enabled.
<code>ndr</code>	Session will be checked by IPS signature
<code>nds</code>	Session will be checked by IPS anomaly
<code>br</code>	Session is being bridged (TP) mode

## Firewall session setup rate

The number of sessions that can be established in a set period of time is useful information. A session is an end-to-end TCP/IP connection for communication with a limited lifespan. If you record the setup rate during normal operation, when you experience problems you can compare the baseline setup rate to the rate that occurs when you're troubleshooting. This can be a useful step to help you define your problem.

A reduced firewall session setup rate can be the result of a number of things, such as a lack of system resources on the FortiGate or reaching the limit of your session count for your VDOM.

### To view your session setup rate method 1 - CLI

```
FGT# get system performance status
CPU states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU0 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU1 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU2 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU3 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 4050332k total, 530512k used (13%), 3376844k free (83%), 142976k freeable (3%)
Average network usage: 131 / 90 kbps in 1 minute, 26 / 15 kbps in 10 minutes, 49 / 42
kbps in 30 minutes
Average sessions: 80 sessions in 1 minute, 30 sessions in 10 minutes, 42 sessions in 30
minutes
Average session setup rate: 3 sessions per second in last 1 minute, 0 sessions per
second in last 10 minutes, 0 sessions per second in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 1 days, 2 hours, 45 minutes
```

The information you're looking for is the average sessions section, in the above output. This example shows that there were 80 sessions in 1 minute, or an average of 3 sessions per second. The values for 10 minutes and 30 minutes allow you to take a longer average for a more reliable value if your FortiGate is working at maximum capacity. The smallest FortiGate can have 1,000 sessions established per second across the unit.

Remember that session setup rate is a global command. If you have multiple VDOMs configured with many sessions in each one, the session setup rate per VDOM will be slower than if there are no VDOMs configured.

## Finding object dependencies

An administrator may not be permitted to delete a configuration object if there are other configuration objects that depend on it. This command identifies other objects which depend on, or make reference to, the configuration object in question. If an error is displayed that an object is in use and can't be deleted, this command can help identify the source of the problem.

Additionally, if you have a virtual interface with objects that depend on it, you need to find and remove those dependencies before you delete the interface.

### CLI method

When you run multiple VDOMs, you run this command in the global configuration only and it searches for the named object in both the most recently used global and VDOM configurations:

```
diagnose system checkused <path.object.mkey>
```

For example, to verify which objects a security policy with an ID of 1 refers to, enter the following command:

```
diagnose system checkused firewall.policy.policyid 1
```

To check what's referred to by interface `port1`, enter the following command:

```
diagnose system checkused system.interface.name port1
```

To show all dependencies for an interface, enter the following command:

```
diagnose system checkused system.interface.name <interface name>
```

### Sample output:

```
entry used by table firewall.address:name '10.98.23.23_host'  
entry used by table firewall.address:name 'NAS'  
entry used by table firewall.address:name 'all'  
entry used by table firewall.address:name 'fortinet.com'  
entry used by table firewall.vip:name 'TORRENT_10.0.0.70:6883'  
entry used by table firewall.policy:policyid '21'  
entry used by table firewall.policy:policyid '14'  
entry used by table firewall.policy:policyid '19'
```

In this example, the interface has dependent objects, including four address objects, one VIP, and three security policies.

### GUI method

In the GUI, you can easily check and remove the object dependencies for an interface.

#### To remove interface object dependencies - GUI

1. Go to **Network > Interfaces**.  
The **Ref.** column displays the number of objects that refer to this interface.
2. Select the number in the **Ref.** column for the interface.  
A window listing the dependencies appears.
3. Use these detailed entries to locate and remove object references to this interface.  
The trash can icon changes from gray when you remove all object dependencies.
4. Remove the interface by selecting the check box for the interface, and select **Delete**.

### Flow trace

To trace the flow of packets through the FortiGate, use the following command:

```
diagnose debug flow trace start
```

Follow packet flow by setting a flow filter, using this command:

```
diagnose debug flow {filter | filter6} <option>
```

If your network uses IPv4, enter `filter`. If your network uses IPv6, enter `filter6`.

One of the following variables replaces `<option>`:

Variable	Description
addr	IPv4 or IPv6 address
clear	clear filter
daddr	destination IPv4 or IPv6 address

Variable	Description
dport	destination port
negate	inverse IPv4 or IPv6 filter
port	port
proto	protocol number
saddr	source address
sport	source port
vd	index of virtual domain; -1 matches all



diagnose debug flow output is recorded as event log messages which are then sent to a FortiCloud or a FortiAnalyzer, if connected. Don't run this command longer than necessary, since it generates significant amounts of data.

### To start flow monitoring with a specific number of packets - CLI:

```
diagnose debug flow trace start <N>
```

### To stop flow tracing at any time using - CLI:

```
diagnose debug flow trace stop
```

The following example shows the flow trace for a device with an IP address of 203.160.224.97:

```
diagnose debug enable
diagnose debug flow filter addr 203.160.224.97
diagnose debug flow show function-name enable
diagnose debug flow trace start 100
```

### Flow trace output example - HTTP

To observe the debug flow trace, connect to the web site at the following address (the display may vary slightly):

```
https://www.fortinet.com
```

Comment: SYN packet received:

```
id=20085 trace_id=209 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

SYN sent and a new session is allocated:

```
id=20085 trace_id=209 func=resolve_ip_tuple line=2799
msg="allocate a new session-00000e90"
```

Lookup for next-hop gateway address:

```
id=20085 trace_id=209 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.11.254 via port6"
```

Source NAT, lookup next available port:

```
id=20085 trace_id=209 func=get_new_addr line=1219
msg="find SNAT: IP-192.168.11.59, port-31925"
direction"
```

Matched security policy. Check to see which policy this session matches:

```
id=20085 trace_id=209 func=fw_forward_handler line=317
msg="Allowed by Policy-3: SNAT"
```

Apply source NAT:

```
id=20085 trace_id=209 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

SYN ACK received:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6, 203.160.224.97:80-
>192.168.11.59:31925) from port6."
```

Found existing session ID. Identified as the reply direction:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=210 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Lookup for next-hop gateway address for reply traffic:

```
id=20085 trace_id=210 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.3.221 via port5"
```

ACK received:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, original
direction"
```

Apply source NAT:

```
id=20085 trace_id=211 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from client:



```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
original direction"
```

Apply source NAT:

```
id=20085 trace_id=212 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from server:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
203.160.224.97:80->192.168.11.59:31925) from port6."
```

Match existing session in reply direction:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=213 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

## Flow trace output example - IPsec (policy-based)

```
id=20085 trace_id=1 msg="vd-root received a packet(proto=1, 10.72.55.240:1->10.71.55.10:8)
from internal."
id=20085 trace_id=1 msg="allocate a new session-00001cd3"
id=20085 trace_id=1 msg="find a route: gw-66.236.56.230 via wan1"
id=20085 trace_id=1 msg="Allowed by Policy-2: encrypt"
id=20085 trace_id=1 msg="enter IPsec tunnel-RemotePhase1"
id=20085 trace_id=1 msg="encrypted, and send to 15.215.225.22 with source 66.236.56.226"
id=20085 trace_id=1 msg="send to 66.236.56.230 via intf-wan1"
id=20085 trace_id=2 msg="vd-root received a packet (proto=1, 10.72.55.240:1-1071.55.10:8)
from internal."
id=20085 trace_id=2 msg="Find an existing session, id-00001cd3, original direction"
id=20085 trace_id=2 msg="enter IPsec ="encrypted, and send to 15.215.225.22 with source
66.236.56.226" tunnel-RemotePhase1"
id=20085 trace_id=2 msgid=20085 trace_id=2 msg="send to 66.236.56.230 via intf-wan1"
```

## Packet sniffing and packet capture

When you troubleshoot networks, it helps to look inside the header of the packets. This helps to determine if the packets, route, and destination are all what you expect. Packet capture can also be called a network tap, packet sniffing, or logic analyzing. FortiOS devices can sniff packets using CLI commands or capture packets using the GUI.

Packet sniffing using CLI commands is well-suited for spot checking traffic, but if you have complex filters to enter it can be a lot of work to enter them each time. You can also save the sniffing output. However, you must log to a file and then analyze the file later.

Packet capture in the GUI makes it easy for you to set up multiple filters at once and run only one or two as you need to. You can also use controls to start and stop capturing when you want to. You download packet capture output to your local computer as a \*.pcap file. You must use a third party application, such as Wireshark, to read \*.pcap files. This method is useful to send information to Fortinet support to help resolve an issue.

The following table presents a comparison between the two methods:

Features	Packet sniffing	Packet capture
Command location	CLI	GUI
Third party software required	puTTY to log plaintext output	Wireshark, or similar application, to read *.pcap files
Read output in plain text file	yes	no
Read output as *.pcap file using Wireshark, or similar application	no	yes
Easily configure single quick and simple filter	yes	no
Record packet interface	yes	no
Configure complex sniffer filters on multiple interface	no	yes
sniff IPv6	hard	easy
sniff non-IP packets	no	yes
Filter packets by protocol and/or port	easy	easy
Filter packets by source and/or destination address	easy	easy

## Packet sniffing

If you're running a constant traffic application, such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the FortiGate is silently dropping packets for reasons such as Reverse Path Forwarding (RPF). RPF, also called anti-spoofing, prevents an IP packet from being forwarded if its source IP doesn't belong to a locally attached subnet (local interface) or isn't part of the routing between the FortiGate and another source (static route, RIP, OSPF, BGP). Note that you can disable RPF by turning on asymmetric routing

in the CLI (config system settings, set asymroute enable), but this disables stateful inspection on the FortiGate and causes many features to be turned off.



If you configure virtual IP addresses on your FortiGate, it will use those addresses instead of the physical IP addresses. You'll notice this when you're sniffing packets because all the traffic uses the virtual IP addresses. This is due to the ARP update that's sent out when the VIP address is configured.

Before you start sniffing packets on the CLI, you should prepare to capture the output to a file. A large amount of data may scroll by and you won't be able to see it without first saving it to a file. One method is to use a terminal program like puTTY to connect to the FortiGate CLI. Once the packet sniffing count is reached, you can end the session and analyze the output in the file.

The general form of the internal FortiOS packet sniffer command is:

```
diagnose sniffer packet <interface_name> <'filter'> <verbose> <count> <tsformat>
```

To stop the sniffer, type CTRL+C.

<b>&lt;interface_name&gt;</b>	The name of the interface to sniff, such as "port1" or "internal". This can also be "any" to sniff all interfaces.
<b>&lt;'filter'&gt;</b>	What to look for in the information the sniffer reads. "none" indicates no filtering, and all packets are displayed as the other arguments indicate.  The filter must be inside single quotes (').
<b>&lt;verbose&gt;</b>	The level of verbosity as one of:  <ol style="list-style-type: none"> <li>1 - print header of packets</li> <li>2 - print header and data from IP of packets</li> <li>3 - print header and data from Ethernet of packets</li> <li>4 - print header of packets with interface name</li> </ol>
<b>&lt;count&gt;</b>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run until you stop it with <CTRL+C>.
<b>&lt;tsformat&gt;</b>	The format of timestamp. <ul style="list-style-type: none"> <li>• a: absolute UTC time, yyyy-mm-dd hh:mm:ss.ms</li> <li>• l: absolute LOCAL time, yyyy-mm-dd hh:mm:ss.ms</li> <li>• otherwise: relative to the start of sniffing, ss.ms</li> </ul>

For a simple sniffing example, enter the CLI command `diagnose sniffer packet port1 none 1 3`. This displays the next three packets on the port1 interface using no filtering, and verbose level 1. At this verbosity level, you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets and that 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diagnose sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
```

```
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757  
0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808  
0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface that are traveling between a computer with the host name of “PC1” and a computer with the host name of “PC2”. With verbosity 4 and above, the sniffer trace displays the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type CTRL+C.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and icmp" 4
```

The following CLI command for a sniffer includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for example, PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

## Packet capture

Packet capture tells you what is happening on the network at a low level. This can be very useful for troubleshooting problems, such as:

- Finding missing traffic
- Seeing if sessions are setting up properly
- Locating ARP problems such as broadcast storm sources and causes
- Confirming which address a computer is using on the network, if they have multiple addresses or are on multiple networks
- Confirming routing is working as you expect
- Connecting wireless clients
- Missing PING packets
- A particular type of packet is having problems, such as UDP, which is commonly used for streaming video

If you're running a constant traffic application such as ping, packet capture can tell you if the traffic is reaching the destination, how the port enters and exits the FortiGate, if the ARP resolution is correct, and if the traffic is returning to the source as expected. You can also use packet switching to verify that NAT or other configuration is translating addresses or routing traffic the way that you want it to.

Before you start capturing packets, you need to have a good idea of what you're looking for. Capture is used to confirm or deny your ideas about what's happening on the network. If you try capture without a plan to narrow your search, you can end up with too much data to effectively analyze. On the other hand, you need to capture enough packets to really understand all of the patterns and behavior that you're examining.

To use packet capture, the FortiGate must have a disk. You can enable the `capture-packet` in the firewall policy, using the following CLI commands:

```
config firewall policy  
edit <id>
```

```

    set capture-packet enable
end

```

To configure packet capture filters, go to **Network > Packet Capture**.

When you add a packet capture filter, enter the following information and select **OK**.

<b>Interface</b>	<p>Select the interface to sniff from the drop-down menu.</p> <p>You must select one interface. You can't change the interface without deleting the filter and creating a new one, unlike the other fields.</p>
<b>Max Packets to Save</b>	<p>Enter the number of packets to capture before the filter stops.</p> <p>This number can't be zero. You can halt the capturing before this number is reached.</p>
<b>Enable Filters</b>	Select this option to specify filter fields
<b>Host(s)</b>	<p>Enter the IP address of one or more hosts</p> <p>Separate multiple hosts with commas. To enter a range, use a dash without spaces, for example 172.16.1.5-172.16.1.15, or enter a subnet.</p>
<b>Port(s)</b>	<p>Enter one or more ports to capture on the selected interface.</p> <p>Separate multiple ports with commas. To enter a range, use a dash without spaces, for example 88-90</p>
<b>VLAN(s)</b>	<p>Enter one or more VLANs (if any).</p> <p>Separate multiple VLANs with commas.</p>
<b>Protocol</b>	Enter one or more protocols. Separate multiple protocols with commas. To enter a range, use a dash without spaces, for example 1-6, 17, 21-25.
<b>Include IPv6 Packets</b>	Select this option if you're troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.
<b>Include Non-IP Packets</b>	The protocols in the list are all IP based except for ICMP (ping). To capture non-IP based packets, select this feature. Examples of non-IP packets include IPsec, IGMP, ARP, and ICMP.

If you select a filter, you have the option to start and stop packet capture in the edit window, or download the captured packets. You can also see the filter status and the number of packets captured.

You can select the filter and start capturing packets. When the filter is running, the number of captured packets increases until it reaches the **Max Packet Count** or you stop it. When the filter is running, you can't download the output file.

When the packet capture is complete, you can download the \*.pcap file. You must use a third party application, such as Wireshark, to read \*.pcap files. This tool provides you with extensive analytics and the full contents of the packets that were captured.

To start, stop, or resume packet capture, use the symbols on the screen. These symbols are the same as those used for audio or video playback. Hover over the symbol to reveal explanatory text. Similarly, to download the \*.pcap file, use the download symbol on the screen.

## NPU-based interfaces

Many Fortinet products contain network processors, such as NP1, NP2, NP4, and NP6.

When you use the NPU-based interfaces, you can see only the initial session setup using the `diagnose debug flow` command. If the session is correctly programmed into the ASIC (fastpath), the `diagnose debug flow` command won't detect the packets that arrive at the CPU. If the NPU functionality is disabled, the CPU detects all packets. However, you should only use this for troubleshooting purposes.

First, obtain the NP4 or NP6 ID and the port numbers, using the following command:

```
diagnose npu {np4|npu6}list
```

### Sample output:

```
ID Model Slot Interface
0 On-board port1 fabric1 fabric3 fabric5
1 On-board fabric2 port2 base2 fabric4
```

Run the following commands:

```
diagnose npu {np4|npu6}fastpaf th disable <dev_id>
```

(where `dev_id` is the NP4 or NP6 number)

Then, run this command:

```
diagnose npu {np4|npu6}fastpath-sniffer enable port1
```

### Sample output:

```
NP4 Fast Path Sniffer on port1 enabled
```

This causes all traffic on **port1** of the network processor to be sent to the CPU. This means that you can take a standard sniffer trace and use other diagnose commands, if it's a standard CPU-driven port.

These commands only apply to the newer NP4 and NP6 interfaces.

## Debug command

Debugging output provides continuous, real-time event information. Debugging output continues until you stop or reboot the unit. Debugging output can affect system performance and is continually generated even though output might not be displayed in the CLI console.

Debugging information that's displayed in the console scrolls in the console display and may prevent you from entering CLI commands, such as, the command to disable the debug display. To turn off debugging output as the display scrolls by, press the **↑** key to recall the recent `diagnose debug` command, press backspace, and type "0", and **Enter**.

To enable debug output display, use the following command:

```
diagnose debug enable
```

Once you enable debug output display, specify the debuggin information that you require, using the following command:

```
diagnose debug <option> <level>
```

Debug command options include the following:

enable	Enable debug output
disable	Disable debug output
info	Show active debug level settings
reset	Reset all debug level to default
report	Report for tech support
crashlog	Crash log info
config-error-log	Configure error log info
sql-log-error	SQL log database error info
application	application
kernel	kernel
remote-extender	remote-extender
cli	Debug CLI
cmdb-trace	Trace CLI
rating	Display rating info
authd	Authentication daemon
fsso-polling	FSSO active directory poll module
flow	Trace packet flow in kernel
urlfilter	urlfilter
admin	Admin user

You can set the debug level at the end of the command. For example, typical values are 2 and 3:

```
diagnose debug application DHCPDS 2
diagnose debug application spamfilter 2
```

Fortinet Support will advise you about which debugging level you should use.

Timestamps can be enabled to the debug output using the following command:

```
diagnose debug console timestamp enable
```

When you finish examining the debug output, disable it, using the following command:

```
diagnose debug disable
```

## Debug output example

This example shows the IKE negotiation for a secure logging connection from a FortiGate to a FortiAnalyzer.

```
diagnose debug reset
diagnose vpn ike log-filter src-addr4 192.168.11.2
diagnose debug enable
```

## Sample output:

```
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2->192.168.10.201:500, natt_mode=0 rekey=0
phase2=FGh_FtiLog1
FGh_FtiLog1: using existing connection, dpd_fail=0
FGh_FtiLog1: found phase2 FGh_FtiLog1
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2 -> 192.168.10.201:500 negotiating
FGh_FtiLog1: overriding selector 225.30.5.8 with 192.168.11.2
FGh_FtiLog1: initiator quick-mode set pfs=1536...
FGh_FtiLog1: try to negotiate with 1800 life seconds.
FGh_FtiLog1: initiate an SA with selectors: 192.168.11.2/0.0.0.0->192.168.10.201,
ports=0/0, protocol=0/0
Send IKE Packet(quick_outI1):192.168.11.2:500(if0) -> 192.168.10.201:500, len=348
Initiator: sent 192.168.10.201 quick mode message #1 (OK)
FGh_FtiLog1: set retransmit: st=168, timeout=6.
```

In this example:

192.168.11.2->192.168.10.201:500	Source and destination gateway IP address
dpd_fail=0	Found existing Phase 1
pfs=1536...	Create new Phase 2 tunnel

## The execute tac report command

The `execute tac report` command runs an exhaustive series of diagnostic commands. It runs commands that are only needed if you're using certain features like high availability (HA), VPN tunnels, or a modem. The report takes a few minutes to finish because of the amount of output that's generated. If you're logging CLI output logged to a file, you can run this command to familiarize yourself with the diagnostic commands.

When you call Fortinet Customer Support, you may be asked to use the output from this CLI command to provide information about your FortiGate and its current state.

## Other commands

### ARP table

To view the ARP cache, use the following command:



```
get system arp
```

To view the ARP cache in the system, use the following command:

```
diagnose ip arp list
```

### Sample output:

```
index=14 ifname=internal 224.0.0.5 01:00:5e:00:00:05 state=00000040 use=72203
confirm=78203 update=72203 ref=1
index=13 ifname=dmz 192.168.3.100 state=00000020 use=1843 confirm=650179 update=644179
ref=2 ? VIP
index=13 ifname=dmz 192.168.3.109 02:09:0f:78:69:ff state=00000004 use=71743 confirm=75743
update=75743 ref=1
index=14 ifname=internal 192.168.11.56 00:1c:23:10:f8:20 state=00000004 use=10532
confirm=10532 update=12658 ref=4
```

To remove the ARP cache, use the following command:

```
execute clear system arp table
```

To remove a single ARP entry, use the following command:

```
diagnose ip arp delete <interface name> <IP address>
```

To add static ARP entries, use the following command:

```
config system arp-table
```

## Time and date settings

Check time and date settings for log message timestamp synchronization (the Fortinet support group may request this) and for certificates that have a time requirement to check for validity. Use the following commands:

```
execute time
current time is: 12:40:48
last ntp sync:Thu Mar 16 12:00:21 2006
execute date
current date is: 2006-03-16
```

To force synchronization with an NTP server, use the following command:

```
config system ntp
set ntpsync {enable|disable}
end
```

If all devices have the same time, it helps to correlate log entries from different devices.

## IP address

There may be times when you want to verify that the IP addresses assigned to the FortiGate interfaces are what you expect them to be. This is easily accomplished from the CLI, using the following command:

```
diagnose ip address list
```

The output from this command lists the IP address and mask (if available), the `index` of the interface (a type of ID number), and the `devname` (the interface name). While physical interface names are set, virtual interface

names can vary. A good way to use this command is to list all of the virtual interface names. For `vsys_ha` and `vsys_fgfm`, the IP addresses are the local host, which are virtual interfaces that are used internally.

```
# diagnose ip address list
IP=10.31.101.100->10.31.101.100/255.255.255.0 index=3 devname=internal
IP=172.20.120.122->172.20.120.122/255.255.255.0 index=5 devname=wan1
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=8 devname=root
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=11 devname=vsys_ha
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=13 devname=vsys_fgfm
```

## FortiOS ports

There are 65,535 ports in TCP and UDP stacks that applications can use when they communicate with each other. Many of these ports are commonly known to be associated with specific applications or protocols. These ports can be useful when you troubleshoot your network.

Use the following ports when you troubleshoot your FortiGate:

Port	Functionality
UDP 53	DNS lookup, RBL lookup
UDP 53 or UDP 8888	FortiGuard Antispam or Web Filtering rating lookup
UDP 53 (default) or UDP 8888 and UDP 1027 or UDP 1031	FDN server list - source and destination port numbers vary by originating or reply traffic
UDP 123	NTP synchronization
UDP 162	SNMP traps
UDP 514	SYSLOG - All FortiOS versions can use syslog to send log messages to remote syslog servers
TCP 22	Configuration backup to FortiManager unit or FortiGuard Analysis and Management Service
TCP 25	SMTP alert email, encrypted virus sample auto-submit
TCP 389 or TCP 636	LDAP or PKI authentication
TCP 443	FortiGuard Antivirus or IPS update - When you request updates from a FortiManager, instead of directly from the FDN, you must reconfigure this port must be as TCP 8890
TCP 443	FortiGuard Analysis and Management Service
TCP 514	FortiGuard Analysis and Management Service log transmission (OFTP)
TCP 514	SSL Management Tunnel to FortiGuard Analysis and Management Service

Port	Functionality
TCP 514	Quarantine, remote access to logs and reports on a FortiAnalyzer unit, device registration with FortiAnalyzer units (OFTP)
TCP 1812	RADIUS authentication
TCP 8000 and TCP 8002	FSSO
TCP 10151	FortiGuard Analysis and Management Service contract validation

## FortiAnalyzer and FortiManager ports

If you have a FortiAnalyzer or FortiManager on your network, you may need to use the following ports to troubleshoot network traffic:

Port	Functionality
UDP 53	DNS lookup
UDP 123	NTP synchronization
UDP 137-138	Windows share
UDP 162	SNMP traps
UDP 514	Syslog, log forwarding
TCP 21 or TCP 22	Log and report upload
TCP 25	SMTP alert email
TCP 389 or TCP 636	User name LDAP queries for reports
TCP 443	RVS update
TCP 1812	RADIUS authentication
TCP 3000	Log aggregation client

## FortiGuard troubleshooting

The FortiGuard service provides updates to AntiVirus (AV), Antispam (AS), Intrusion Protection Services (IPS), Webfiltering (WF), and more. The FortiGuard Distribution System (FDS) consists of a number of servers across the world that provide updates to your FortiGate. Problems can occur with connection to FDS, and its

configuration on your local FortiGate. Some of the more common troubleshooting methods are listed here, including:

- [Troubleshooting process for FortiGuard updates](#)
- [FortiGuard server settings](#)

## Troubleshooting process for FortiGuard updates

The following process shows the logical steps that you should take when you troubleshoot problems with FortiGuard updates:

**1. Does the device have a valid licence that includes these services?**

Each device requires a valid FortiGuard license to access updates for some or all of these services. You can verify the status of the support contract for your devices at the [Fortinet Support website](#).

**2. If the device is part of a "high availability (HA)" cluster, do all members of the cluster have the same level of support?**

As with the previous step, you can verify the support contract status for all of the devices in your HA cluster at the Fortinet Support website.

**3. Are services enabled on the device?**

To see the FortiGuard information and status for a device, in the GUI go to **System > FortiGuard**. On that page you can verify the status of each component, and enable each service (if required). If there are problems, see the [FortiGuard](#) discussion in the [Fortinet Communication Ports and Protocols](#) chapter of the FortiOS Handbook.

**4. Can the device communicate with FortiGuard servers?**

Go to **System > FortiGuard** in the GUI and try to update AV and IPS, or test the availability of WF and AS default and alternate ports. If you encounter problems, see the [FortiGuard](#) discussion in the [Fortinet Communication Ports and Protocols](#) chapter of the FortiOS Handbook.

**5. Is there proper routing to reach the FortiGuard servers?**

Ensure there is a static or dynamic route that allows your FortiGate to reach the FortiGuard servers. Usually a generic default route to the internet is enough, but you may need to verify this if your network is complex.

**6. Are there issues with DNS?**

An easy way to test this is to attempt a traceroute from behind the FortiGate to an external network using the Fully Qualified Domain Name (FQDN) for a location. If the traceroute FQDN name doesn't resolve, you have general DNS problems.

**7. Is there anything upstream that might be blocking FortiGuard traffic, either on the network or ISP side?**

Many firewalls block all ports, by default, and ISPs often block ports that are low. There may be a firewall between the FortiGate and the FortiGuard servers that's blocking the traffic. FortiGuard uses port 53, by default, so if it's being blocked you need to either open a hole for it or change the port it's using.

**8. Is there an issue with source ports?**

It's possible that ports that the FortiGate uses to contact FortiGuard are being changed before they reach FortiGuard or on the return trip before they reach your FortiGate. A possible solution for this is to use a fixed-port at NAT'd firewalls to ensure the port remains the same. You can use packet sniffing to find more information about what's happening with ports.

**9. Are there security policies that include antivirus?**

If none of the security policies include antivirus, the antivirus database won't be updated. If antivirus is included, only the database type that's used will be updated.

## FortiGuard server settings

Your local FortiGate connects to remote FortiGuard servers to get updates to FortiGuard information, such as new viruses that may have been found or other new threats. This section shows ways that you can to display FortiGuard server information on your FortiGate, and how you can use that information and update it to fix potential problems.

### Displaying the server list

The `get webfilter status` or `diagnose debug rating` command shows the list of FDS servers that the FortiGate uses to send web filtering requests. Rating requests are only sent to the server at the top of the list in normal operation. Each server is probed for Round Trip Time (RTT) every two minutes.

Optionally, you can add a refresh rate to the end of this command to determine how often the server list is refreshed.

Rating may not be enabled on your FortiGate.

#### To show the list of servers a FortiGate uses to send web filtering requests - CLI

```
get webfilter status
```

Sample output:

```
Locale : english
License : Contract
Expiration : Thu Oct 9 02:00:00 2011
--- Server List (Mon Feb 18 12:55:48 2008) ---
```

IP	Weight	RTT	Flags	TZ	Packets	CurrLost	TotalLost
a.b.c.d	0	1	DI	2	1926879	0	11176
10.1.101.1	10	329		1	10263	0	633
10.2.102.2	20	169		0	16105	0	80
10.3.103.3	20	182		0	6741	0	776
10.4.104.4	20	184		0	5249	0	987
10.5.105.5	25	181		0	12072	0	178

### Output details

The server list includes the IP addresses of alternate servers if the first entry can't be reached. In this example, the IP addresses are not public addresses.

The following flags in `get webfilter status` indicate the server status:

Flag	Description
D	<p>The server was found through the DNS lookup of the hostname.</p> <p>If the hostname returns more than one IP address, all of them are flagged with D and are used first for INIT requests before falling back to the other servers.</p>
I	The server to which the last INIT request was sent
F	The server hasn't responded to requests and is considered to have failed
T	The server is currently being timed
S	<p>Rating requests can be sent to the server</p> <p>The flag is set for a server only in two cases:</p> <ol style="list-style-type: none"> <li>1. The server exists in the servers list received from the FortiManager or any other INIT server.</li> <li>2. The server list received from the FortiManager is empty so the FortiManager is the only server that the FortiGate knows, and it should be used as the rating server.</li> </ol>

### Sorting the server list

The server list is sorted first by weight. The server with the smallest RTT appears at the top of the list, regardless of weight. When a packet is lost (there has been no response in 2 seconds), it's re-sent to the next server in the list. Therefore, the top position in the list is selected based on RTT, while the other positions are based on weight.

### Calculating weight

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a remote server, the weight isn't allowed to dip below a base weight. The base weight is calculated as the difference in hours between the FortiGate and the server multiplied by 10. The farther away the server is, the higher its base weight is and the lower it appears in the list.

# Troubleshooting tips

The following tips present common causes of problems.

## How to check hardware connections

- Are all of the cables and interfaces connected properly?
- Is the LED for the interface green?

## How to check FortiOS network settings

- If you're having problems connecting to the management interface, is your protocol enabled on the interface for administrative access?
- Does the interface have an IP address?

## How to check CPU and memory resources

- Is the FortiGate's CPU running at almost 100 percent usage?
- Is your FortiGate running low on memory?

## How to check modem status

- Is the modem connected?
- Are there PPP issues?

## How to run ping and traceroute

- Is the FortiGate experiencing complete packet loss?

## How to check the logs

- Do you need to identify a problem?

## How to verify the contents of the routing table (in NAT mode)

- Are there routes in the routing table for default and static routes?
- Do all connected subnets have a route in the routing table?
- Does a route have a higher priority than it should?

## How to verify the correct route is being used

- Is the traffic routed correctly?

## How to verify the correct firewall policy is being used

- Is the correct firewall policy applied to the expected traffic?

## How to check the bridging information in transparent mode

- Are you having problems in transparent mode?

## How to check the number of sessions that UTM proxy uses

- Have you reached the maximum number of sessions for a protocol?
- Are new sessions failing to start for a certain protocol?

## How to examine the firewall session list

- Are there active firewall sessions?

#### How to check wireless information

- Is the wireless network working properly?

#### How to verify connectivity to FortiGuard

- Is the FortiGate communicating properly with FortiGuard?

#### How to perform a sniffer trace (CLI and packet capture)

- Is traffic entering the FortiGate? Does the traffic arrive on the expected interface?
- Is the ARP resolution correct for the next-hop destination?
- Is the traffic exiting the FortiGate to the destination as expected?
- Is the FortiGate sending traffic back to the originator?

#### How to debug the packet flow

- Is traffic entering or leaving the FortiGate as expected?

## How to check hardware connections

If there's no traffic flowing from the FortiGate, you may have a hardware problem.

To check hardware connections:

- Ensure network cables are plugged into the interfaces.
- Verify that the LED connection lights for the network cables are the right color (usually green).
- If the cable or its connector are damaged, change the cable. You should also change the cable if you're not sure about the type or quality of the cable, such as straight through or crossover, or if you see exposed wires at the connector.
- Connect the FortiGate to different hardware.
- Ensure the link status is set to **Up** for the interface (see **Network > Interfaces**). The link status is based on the physical connection and can't be set in FortiOS.

If any of these solve the problem, it was a hardware connection problem. You should still perform some basic software connectivity tests to ensure complete connectivity. The interface might also be disabled, or its **Administrative Status** is set to **Down**.

#### To enable an interface - GUI

1. Go to **Network > Interfaces**.
2. Select and edit the interface to enable, such as **port1**.
3. Find **Administrative Status** at the bottom of the screen, and select **Up**.
4. Select **Apply**.

#### To enable an interface - CLI

```
config system interface
  edit port1
    set status up
  next
end
```



## How to check FortiOS network settings

You can manage FortiOS network settings in the GUI and the CLI. The following information includes troubleshooting and best practice information. The network settings include:

- [Interface settings](#)
- [DNS settings](#)
- [DHCP server settings](#)

### Interface settings

If you can access the FortiGate with the management cable only, the first step is to display the interface settings. To display the settings for the internal interface, use the following CLI command:

```
FGT# show system interface <interface_name>
```

For a complete listing of all the possible interface settings, use the following CLI command:

```
config system interface
  edit <interface_name>
    get
  end
```

Check the interface settings to ensure that they aren't preventing traffic. Check the following items (only the GUI terms are shown, CLI terms may vary):

Setting	Description
<b>Link Status</b>	<b>Down</b> until a valid cable is plugged into this interface, then it will be <b>Up</b> .  The Link Status is shown physically by the connection LED for the interface. If the LED is green, the connection is good. If Link Status is <b>Down</b> , the interface doesn't work. Link Status is also displayed on the <b>Network &gt; Interfaces</b> screen, by default.
<b>Addressing mode</b>	Don't use <b>DHCP</b> if you don't have a DHCP server. You won't be able to log in to an interface in DHCP mode as it won't have an IP address
<b>IP/Network Mask</b>	An interface needs an IP address to be able to connect to other devices. Ensure there is a valid IP address in this field. The one exception is if <b>DHCP</b> is enabled for this interface to get its IP address from an external DHCP server.
<b>IPv6 address</b>	The same protocol must be used by both ends to complete the connection. Ensure both this interface and the remote connection are both using IPv4 or both using IPv6 addressing.
<b>Administrative access</b>	If no protocols are selected, you will have to use the local management cable to connect to the unit. If you're using IPv6, configure the IPv6 administrative access protocols.
<b>Administrative status</b>	Set to <b>Up</b> or the interface won't work.

## DNS settings

You can trace many networking problems back to DNS issues. Check the following items:

- Are there values for both primary and secondary entries?
- Is the local domain name correct?
- Are you using IPv6 addressing? If so, are the IPv6 DNS settings correct?
- Are you using Dynamic DNS (DDNS)? If so, is it using the correct server, credentials, and interface?
- Can you contact both DNS servers to verify the servers are operational?
- If an interface addressing mode is set to DHCP and is set to override the internal DNS, is that interface receiving a valid DNS entry from the DHCP server? Is it a reasonable address and can it be contacted to verify it's operational?
- Are there any DENY security policies that need to allow DNS?
- Can any internal device perform a successful traceroute to a location using the FQDN?

## DHCP server settings

DHCP servers are common on internal and wireless networks. If the DHCP server isn't configured correctly, it can cause problems. Check the following items:

- Is the DHCP server entry set to **Relay**? If so, verify there is another DHCP server to which requests can be relayed. Otherwise, it should be set to **Server**.
- Is the DHCP server enabled?
- Does the DHCP server use a valid IP address range? Are other devices using the addresses? If one or more devices are using IP addresses in this range, you can use the IP reservation feature to ensure the DHCP server doesn't use these addresses.
- Is there a gateway entry? If not, add a gateway entry to ensure that the server's clients have a default route.
- Is the system DNS setting being used? A best practice is to avoid confusion by using the system DNS whenever possible. However, you can specify up to three custom DNS servers, and you should use all three entries for redundancy.



There are some situations, such as a new wireless interface, or during the initial FortiGate configuration, where interfaces override the system DNS entries. When this happens, it often shows up as intermittent Internet connectivity. To fix the problem, go to **Network > DNS** and enable **Use FortiGuard Servers**.

## How to check CPU and memory resources

System resources are shared and a number of processes run simultaneously on the FortiGate unit.

The **Resource** widgets for **CPU** and **Memory** on the **Dashboard** provide a quick way to monitor usage.

To use the CLI to check the system resources on your FortiGate unit, run the following command:

```
FGT# get system performance status
```

This command provides a quick and easy snapshot of the FortiGate.

The first line of output shows the CPU usage by category. A FortiGate that is doing nothing looks like the following example:

```
CPU states: 0% user 0% system 0% nice 100% idle
```

However, if your network is running slowly, you might see something like the following example:

```
CPU states: 1% user 98% system 0% nice 1% idle
```

This line shows that all of the CPU is used up by system processes. Normally this shouldn't happen since it shows that the FortiGate is overloaded for some reason. If you see this overloading, you should investigate further because it's possible that a process, such as `scanunitid`, is using all of the resources to scan traffic. In this case, you must reduce the amount of traffic that's being scanned by blocking unwanted protocols, configuring more security policies to limit scanning to certain protocols, or similar actions. It's also possible that a hacker has gained access to your network and is overloading it with malicious activity, such as running a spam server or using zombie PCs to attack other networks on the Internet. You can use the `get system performance top` CLI command to get more information about the CPU. This command shows you all of the top processes that are running on the FortiGate (the names are on the left) and their CPU usage. If a process is using most of the CPU cycles, investigate it to determine whether the activity is normal.

The second line of output from the `get system performance status` command shows the memory usage. Memory usage shouldn't exceed 90%. If memory is too full, some processes won't be able to function properly. For example, if the system is running low on memory, antivirus scanning enters into failopen mode where it drops connections or bypasses the antivirus system.

The other lines of output, such as average network usage, average session setup rate, viruses caught, and IPS attacks blocked, can also help you determine why system resource usage is high. For example, if network usage is high, it results in high traffic processing on the FortiGate, or if the session setup rate is very low (or zero) the proxy may be overloaded and unable to do its job.

## How to troubleshoot high memory usage

As with any system, a FortiGate has limited hardware resources, such as memory, and all processes running on the FortiGate share the memory. Each process uses more or less memory, depending on its workload. For example, a process usually uses more memory in high traffic situations. If some processes use all of the available memory, other processes won't be able to run.

When high memory usage occurs, the services may freeze up, connections may be lost, or new connections may be refused.

If you see high memory usage in the **Memory** widget, the FortiGate may be handling high traffic volumes. Alternatively, the FortiGate may have connection pool limits that are affecting a single proxy. If the FortiGate receives a large volume of traffic on a specific proxy, the unit may exceed the connection pool limit. If the number of free connections within a proxy connection pool reaches zero, issues may occur.

Use the following CLI command, which uses the antivirus failopen feature:

```
config system global
    set av-failopen idledrop
end
```

If you set `av-failopen` to `idledrop`, the FortiGate drops connections based on the client that has the most open connections. This helps you determine the behavior of the FortiGate antivirus system if it becomes overloaded in high traffic.

Use the following CLI command, which shows information about current memory usage:

```
diagnose hardware sysinfo memory
```

### Sample output:

```
total: used: free: shared: buffers: cached: shm:
Mem: 2074185728 756936704 1317249024 0 20701184 194555904 161046528
Swap:      0      0      0
MemTotal:   2025572 kB
MemFree:    1286376 kB
MemShared:      0 kB
Buffers:     20216 kB
Cached:      189996 kB
SwapCached:    0 kB
Active:      56644 kB
Inactive:    153648 kB
HighTotal:      0 kB
HighFree:      0 kB
LowTotal:     2025572 kB
LowFree:     1286376 kB
SwapTotal:      0 kB
SwapFree:      0 kB
```

## How to troubleshoot high CPU usage

If you deploy too many FortiOS features at the same time, you can overextend the CPU resources on a FortiGate. When this occurs, the FortiGate experiences connection-related problems.

Some examples of CPU intensive features are:

- VPN high-level encryption
- Intensive scanning of all traffic
- Logging all traffic and packets
- Dashboard widgets that frequently perform data updates

### 1. Determine the current level of CPU usage.

There are two ways to do this. The simplest is to look at the **CPU** widget on the **Dashboard**. The real-time CPU usage is displayed for different timeframes. The second method provides precise usage values both for overall usage and for specific processes. To use it, run the `diagnose system top` command in the CLI.

### Sample output:

```
Run Time: 86 days, 0 hours and 10 minutes
0U, 0N, 0S, 100I, 0WA, 0HI, 0SI, 0ST; 3040T, 2437F
bcm.user 93 S < 3.1 0.4
httpsd 18922 S 1.5 0.5
httpsd 19150 S 0.3 0.5
newcli 20195 R 0.1 0.1
cmdbsvr 115 S 0.0 0.8
pyfcgid 20107 S 0.0 0.6
forticron 146 S 0.0 0.5
httpsd 139 S 0.0 0.5
cw_acd 166 S 0.0 0.5
miglogd 136 S 0.0 0.5
```

```

pyfcbid 20110 S 0.0 0.4
pyfcbid 20111 S 0.0 0.4
pyfcbid 20109 S 0.0 0.4
httpsd 20192 S 0.0 0.4
miglogd 174 S 0.0 0.4
miglogd 175 S 0.0 0.4
fgfmd 165 S 0.0 0.3
newcli 20191 S 0.0 0.3
initXXXXXXXXXX 1 S 0.0 0.3
httpsd 184 s 0.0 0.3

```

Where the codes on the second output line mean the following:

- **U** is the percentage of user space applications that are currently using the CPU.
- **N** is the percentage of time that the CPU spent on low priority processes since the last shutdown.
- **S** is the percentage of system processes (or kernel processes) that are using the CPU.
- **I** is the percentage of idle CPU resources.
- **WA** is the percentage of time that the CPU spent waiting on IO peripherals since the last shutdown.
- **HI** is the percentage of time that the CPU spent handling hardware interrupt routines since the last shutdown.
- **SI** is the percentage of time that the CPU spent handling software interrupt routines since the last shutdown.
- **ST**, or steal time, is the percentage of time that a virtual CPU waits for the physical CPU when the hypervisor is servicing another virtual processor.
- **T** is the total FortiOS system memory, in MB.
- **F** is free memory, in MB.

Each additional line of the command output displays information for specific processes running on the FortiGate unit. For example, the third line of the output is:

```
newcli 20195 R 0.1 0.1
```

Where:

- **newcli** is the process name. Other process names can include `ipsengine`, `sshd`, `cmdbsrv`, `httpsd`, `scanunitd`, and `miglogd`. Duplicate process names indicate that separate instances of that process are running.
- **20195** is the process ID, which can be any number.
- **R** is the current state of the process. The process state can be:
  - **R** - running
  - **S** - sleep
  - **Z** - zombie
  - **D** - disk sleep
- **0.1** is the percentage of CPU capacity that the process is using. CPU usage can range from 0.0 for a process that is sleeping to higher values for a process that's taking a lot of CPU time.
- **0.1** is the amount of memory that the process is using. Memory usage can range from 0.1 to 5.5 and higher.

When `diagnose system top` is running, enter the following single-key commands :

- Press **q** to quit and return to the normal CLI prompt.
- Press **p** to sort the processes by the amount of CPU resources that the processes are using.
- Press **m** to sort the processes by the amount of memory that the processes are using.

The processes listed are the top processes that are running, and not all of the processes. For example, if 20 processes are listed, they are the top 20 processes, sorted by either CPU or memory usage. You can configure the number of processes displayed, using the following CLI command:

```
diagnose system top <integer_seconds> <integer_maximum_lines>
```

Where:

- `<integer_seconds>` is the delay in seconds (default is 5)
- `<integer_maximum_lines>` is the maximum number of lines (or processes) to list (default is 20).

## 2. Determine what features are using most of the CPU resources.

There is a command in the CLI that shows the top few processes that are currently running and use the most CPU resources. The `get system performance top` CLI command shows a table of information. The second column from the right shows CPU usage by percentage. If the top few entries are using most of the CPU, note which processes they are and try and reduce their CPU load.

Some examples of processes you'll see include:

- `ipsengine` — the IPS engine that scans traffic for intrusions
- `scanunitd` — antivirus scanner
- `httpsd` — secure HTTP
- `iked` — internet key exchange (IKE) in use with IPsec VPN tunnels
- `newcli` — active whenever you're accessing the CLI
- `sshd` — there are active secure socket connections
- `cmdbsrv` — the command database server application

Go to the features that are at the top of the list and look for evidence of them overusing the CPU. Generally, the monitor for a feature is a good place to start.

## 3. Check for unnecessary CPU “wasters”.

These are some best practices that will reduce your CPU usage, even if the FortiGate isn't experiencing high CPU usage.

- Use hardware acceleration wherever possible to offload tasks from the CPU. Offloading tasks, such as encryption, frees up the CPU for other tasks.
- Avoid the use of GUI widgets that require computing cycles, such as the Top Sessions widget. These widgets constantly poll the system for information, which uses CPU and other resources.
- Schedule antivirus, IPS, and firmware updates during off peak hours. These updates don't usually consume CPU resources but they can disrupt normal operation.
- Check the log levels and which events are being logged. This is the severity of the messages that are recorded. Consider going up one level to reduce the amount of logging. Also, if there are events you don't need to monitor, remove them from the list.
- Log to FortiCloud instead of logging to memory or disk. Logging to memory quickly uses up resources and logging to the local disk impacts overall performance and reduces the lifetime of the unit. Fortinet recommends that you log to FortiCloud because it doesn't use much CPU.
- If the disk is almost full, transfer the logs or data off the disk to free up space. When a disk is almost full it consumes a lot of resources to find the free space and organize the files.
- If the FortiGate has packet logging enabled, consider disabling it. When packet logging is enabled, it records every packet that comes through that policy.
- Halt all sniffers and traces.

- Ensure that the FortiGate isn't scanning traffic twice. If traffic enters the FortiGate on one interface, goes out another, and then comes back in again, that traffic doesn't need to be rescanned. Doing so is a waste of resources. However, ensure that traffic is being scanned once.
- Reduce the session timers to close unused sessions faster. Enter the following CLI commands, which reduce the default values. These values reduce the values from defaults. Note that, by default, the system adds 10 seconds to `tcp-timewait`.

```
config system global
    set tcp-halfclose-timer 30
    set tcp-halfopen-timer 30
    set tcp-timewait-timer 0
    set udp-idle-timer 60
end
```

- In **System > Feature Visibility**, only enable features that you need.

4. When CPU usage is under control, use SNMP to monitor CPU usage. Alternatively, use logging to record CPU and memory usage every 5 minutes.

Once things are back to normal, you should set up a warning system that sends you alerts when CPU resources are used excessively. A common method to do this is using SNMP. SNMP monitors many values in FortiOS and allows you to set high water marks that generate events. You run an application on your computer to watch for and record these events. Go to **System > SNMP** to enable and configure an SNMP community. If this method is too complicated, you can use the **System Resources** widget to record CPU usage. However, this method only records problems as they happen and won't send you alerts for problems.

## How to check modem status

If the modem doesn't work properly, or a FortiGate doesn't detect the modem, you can use the following diagnostic command to help you troubleshoot issues with the modem:

```
diagnose system modem {cmd | com | detect | history | external-modem | query| reset}
```

You should always run the following diagnose command after you insert the USB modem into the FortiGate:

```
diagnose system modem detect
```

You can view the modem configuration, using the `get system modem` command. You can also view the modem's vendor and the custom product identification number from the information output that the `get system modem` command shows.

If there are connectivity issues, use the following to help you resolve them:

- `diagnose debug enable` – activates the debug on the console
- `diagnose debug application modemd` – dumps communication between the modem and the unit.
- `diagnose debug application ppp` – dumps the PPP negotiating messages.
- `execute modem dial` – displays modem debug output.

The modem diagnose output shouldn't contain errors on the way to initializing. You should also verify the number that is used to dial into your ISP.

## How to run ping and traceroute

Ping and traceroute are useful tools in network troubleshooting. Alone, either one can determine network connectivity between two points. However, ping can be used to generate simple network traffic that you can view using diagnose commands on the FortiGate. This combination can be very powerful when you're trying to locate network problems.

In addition to their normal uses, ping and traceroute can tell you if your computer or network device has access to a domain name server (DNS). While both tools can use IP addresses alone, they can also use domain names for devices. This is an added troubleshooting feature that can be useful in determining why particular services, such as email or web browsing, may not work properly.



If ping doesn't work, it may be disabled on at least one of the interface settings, and security policies for that interface.

Both ping and traceroute require particular ports to be open on firewalls, or else they can't function. Since you typically use these tools to troubleshoot, you can allow them in the security policies and on interfaces only when you need them. Otherwise, keep the ports disabled for added security.

## Ping

The ping command sends a very small packet to a destination, and waits for a response. The response has a timer that expires when the destination is unreachable.

Ping is part of Layer 3 on the OSI Networking Model. Ping sends Internet Control Message Protocol (ICMP) "echo request" packets to the destination, and listens for "echo response" packets in reply. However, many public networks block ICMP packets because ping can be used in a denial of service (DoS) attack (such as Ping of Death or a smurf attack), or by an attacker to find active locations on the network. By default, FortiGate units have ping enabled while broadcast-forward is disabled on the external interface.

### What ping can tell you

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is some packet loss detected, you should investigate the following:

- Possible ECMP, split horizon, or network loops.
- Cabling, to ensure no loose connections.
- Verify which security policy was used (use the packet count column on the **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6 Policy** page).

If there is total packet loss, you should investigate the following:

- **Hardware:** ensure cabling is correct, and all equipment between the two locations is accounted for.
- **Addresses and routes:** ensure all IP addresses and routing information along the route is configured as expected.
- **Firewalls:** ensure all firewalls, including FortiGate security policies allow PING to pass through.



## How to use ping

Ping syntax is the same for nearly every type of system on a network.

### To ping from a FortiGate

1. Connect to the CLI either through telnet or through the CLI widget on the **Dashboard**.
2. Enter `execute ping 10.11.101.101` to send 5 ping packets to the destination IP address. There are no options for this command.

#### Sample output:

```
Head_Office_620b # execute ping 10.11.101.101
PING 10.11.101.101 (10.11.101.101): 56 data bytes
64 bytes from 10.11.101.101: icmp_seq=0 ttl=255 time=0.3 ms
64 bytes from 10.11.101.101: icmp_seq=1 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=2 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=3 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=4 ttl=255 time=0.2 ms

--- 10.11.101.101 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.2/0.2/0.3 ms
```

### To ping from an Microsoft Windows PC

1. Open a command window.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate with four packets.

Other options include:

- `-t` to send packets until you press "Ctrl+C"
- `-a` to resolve addresses to domain names where possible
- `-n X` to send X ping packets and stop

#### Sample output:

```
C:\>ping 10.11.101.101

Pinging 10.11.101.101 with 32 bytes of data:
Reply from 10.11.101.101: bytes=32 time=10ms TTL=255
Reply from 10.11.101.101: bytes=32 time<1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255

Ping statistics for 10.11.101.101:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 10ms, Average = 3ms
```

### To ping from a Linux PC

1. Go to a shell prompt.
2. Enter `"ping 10.11.101.101"`.

## Traceroute

Where ping will only tell you if it reached its destination and returned successfully, traceroute shows each step of the journey to its destination and how long each step takes. If ping finds an outage between two points, you can use traceroute to locate exactly where the problem is.

### What's traceroute

Traceroute works by sending ICMP packets to test each hop along the route. It sends three packets, and then increases the time to live (TTL) setting by one each time. This effectively allows the packets to go one hop farther along the route. This is the reason why most traceroute commands display their maximum hop count before they start tracing the route, which is the maximum number of steps it takes before it declares the destination unreachable. Also, the TTL setting may result in steps along the route timing out due to slow responses. There are many possible reasons for this to occur.

By default, traceroute uses UDP datagrams with destination ports numbered from 33434 to 33534. The traceroute utility may also offer the option to select use of ICMP echo request (type 8) instead, which the Windows tracert utility uses. If you must, allow both protocols inbound through the FortiGate security policies (UDP with ports from 33434 to 33534 and ICMP type 8).

You can also use the packet count column of the **Policy & Objects > IPv4 Policy** (or **Policy & Objects > IPv6 Policy**, if applicable) page to track traceroute packets. This allows you to verify the connection and confirm which security policy the traceroute packets are using.

### What traceroute can tell you

Ping and traceroute have similar functions, which is to verify connectivity between two points. The big difference is that traceroute shows you each step of the way, where ping doesn't. Also, ping and traceroute use different protocols and ports, so one may succeed where the other fails.

You can verify your DNS connection using traceroute. If you enter an FQDN instead of an IP address for the traceroute, DNS tries to resolve that domain name. If the name isn't resolved, you have DNS issues.

### How to use traceroute

The traceroute command varies slightly between operating systems. Note that in MicroSoft Windows, the command name is shortened to "tracert". Also, your output lists different domain names and IP addresses along your route.

#### To use traceroute on an MicroSoft Windows PC

1. Open a command window.
2. Enter "tracert fortinet.com" to trace the route from the PC to the Fortinet web site.

#### Sample output:

```
C:\>tracert fortinet.com

Tracing route to fortinet.com [208.70.202.225]
over a maximum of 30 hops:
  1  <1 ms  <1 ms  <1 ms  172.20.120.2
  2  66 ms  24 ms  31 ms  209-87-254-xxx.storm.ca [209.87.254.221]
  3  52 ms  22 ms  18 ms  core-2-g0-0-1104.storm.ca [209.87.239.129]
  4  43 ms  36 ms  27 ms  core-3-g0-0-1185.storm.ca [209.87.239.222]
```

```

5 46 ms 21 ms 16 ms te3-x.1156.mpd01.cogentco.com [38.104.158.69]
6 25 ms 45 ms 53 ms te8-7.mpd01.cogentco.com [154.54.27.249]
7 89 ms 70 ms 36 ms te3-x.mpd01.cogentco.com [154.54.6.206]
8 55 ms 77 ms 58 ms sl-st30-chi-.sprintlink.net [144.232.9.69]
9 53 ms 58 ms 46 ms sl-0-3-3-x.sprintlink.net [144.232.19.181]
10 82 ms 90 ms 75 ms sl-x-12-0-1.sprintlink.net [144.232.20.61]
11 122 ms 123 ms 132 ms sl-0-x-0-3.sprintlink.net [144.232.18.150]
12 129 ms 119 ms 139 ms 144.232.20.7
13 172 ms 164 ms 243 ms sl-321313-0.sprintlink.net [144.223.243.58]
14 99 ms 94 ms 93 ms 203.78.181.18
15 108 ms 102 ms 89 ms 203.78.176.2
16 98 ms 95 ms 97 ms 208.70.202.225

```

Trace complete.

The first column on the left shows the hop count, which can't exceed 30 hops. When that number is reached, the traceroute ends.

The second, third, and fourth columns display how much time each of the three packets takes to reach this stage of the route. These values are in milliseconds and normally vary quite a bit. Typically a value of <1ms indicates a local connection.

The fifth column (farthest to the right) shows the domain name of the device and its IP address, or possibly only the IP address.

### To perform a traceroute on a Linux PC

1. Go to a command line prompt.
2. Enter "traceroute fortinet.com".

The Linux traceroute output is very similar to the MicroSoft Windows tracert output.

### To trace a route from a FortiGate to a destination IP address - CLI

```

# execute traceroute www.fortinet.com
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
1 172.20.120.2 0.637 ms 0.653 ms 0.279 ms
2 209.87.254.221 <static-209-87-254-221.storm.ca> 2.448 ms 2.519 ms 2.458 ms
3 209.87.239.129 <core-2-g0-2.storm.ca> 2.917 ms 2.828 ms 9.324 ms
4 209.87.239.199 <core-3-bdi1739.storm.ca> 13.248 ms 12.401 ms 13.009 ms
5 216.66.41.113 <v502.core1.tor1.he.net> 17.181 ms 12.422 ms 12.268 ms
6 184.105.80.9 <100gel-2.core1.nyc4.he.net> 21.355 ms 21.518 ms 21.597 ms
7 198.32.118.41 <ny-paix-gni.twgate.net> 83.297 ms 84.416 ms 83.782 ms
8 203.160.228.217 <217-228-160-203.TWGATE-IP.twgate.net> 82.579 ms 82.187 ms 82.066 ms
9 203.160.228.229 <229-228-160-203.TWGATE-IP.twgate.net> 82.055 ms 82.455 ms 81.808 ms
10 203.78.181.2 82.262 ms 81.572 ms 82.015 ms
11 203.78.186.70 83.283 ms 83.243 ms 83.293 ms
12 66.171.127.177 84.030 ms 84.229 ms 83.550 ms
13 66.171.121.34 <www.fortinet.com> 84.023 ms 83.903 ms 84.032 ms
14 66.171.121.34 <www.fortinet.com> 83.874 ms 84.084 ms 83.810 ms

```

## How to check the logs

You might forget this step in troubleshooting, but it's an important one. Logging records the traffic that passes through the FortiGate to your network and what action the FortiGate takes when it scans the traffic. This information record is called a log message.

When you first configure FortiOS, log as much information as you can. If the logs that the FortiGate generates are too large, you can turn off or scale back the logging for features that you're not using.

As with most troubleshooting steps, before you can determine if the logs indicate a problem, you need to know what logs result from normal operation. Without a baseline it's difficult to troubleshoot.

When you troubleshoot with log files:

- Compare current logs to a recorded baseline of normal operation.
- If you need to increase the level of logging (for example, from Warning to Information) to obtain more information.

When you increase logging levels, ensure that you configure email alerts and select both disk usage and log quota. This ensures that you'll be notified if the increase in logging causes problems. Configure log settings by going to **Log & Report > Log Settings**.

Determine the activities that generate the most log entries:

- Check all logs to ensure important information isn't overlooked.
- Filter or order log entries based on different fields, such as level, service, or IP address, to look for patterns that may indicate a specific problem, such as frequent blocked connections on a specific port for all IP addresses.

Logs help identify and locate any problems, but they won't solve the problems. The purpose of logs is to speed up your problem solving and save you time and effort.

For more information about logging and log reports, see the [Logging and Reporting handbook chapter](#).

## How to verify the contents of the routing table (in NAT mode)

When a FortiGate has limited or no connectivity, a good place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves time and resources that you would spend performing a lookup. If the routing table is full and a new route needs to be added, the oldest, least-used route is removed and the new route is added. This ensures that the most recently used routes stay in the table. If your FortiGate is in transparent mode, you can't perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table, including local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the FortiGate GUI, use the Routing Monitor by going to **Monitor > Routing Monitor**.

In the CLI, use the command `get router info routing-table all`.

### Sample output:

```
FGT# get router info routing-table all

Codes:
  K - kernel, C - connected, S - static, R - RIP, B - BGP
  O - OSPF, IA - OSPF inter area
  N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
  E1 - OSPF external type 1, E2 - OSPF external type 2
  i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
  * - candidate default

S* 0.0.0.0/0 [10/0] via 172.20.120.2, wan1
C 10.31.101.0/24 is directly connected, internal
C 172.20.120.0/24 is directly connected, wan1
```

## How to verify the correct route is being used

If you have more than one default route and want to make sure that traffic is flowing as expected and through the right route, you can run a trace route from a machine in the local area network (LAN). This shows you the first hop that the traffic goes through.

### Sample output:

```
C:\>tracert www.fortinet.com

Tracing route to www.fortinet.com [66.171.121.34]
over a maximum of 30 hops:

 1 <1 ms <1 ms <1 ms 10.10.1.99
 2 1 ms <1 ms <1 ms 172.20.120.2
 3 3 ms 3 ms 3 ms static-209-87-254-221.storm.ca [209.87.254.221]
 4 3 ms 3 ms 3 ms core-2-g0-2.storm.ca [209.87.239.129]
 5 13 ms 13 ms 13 ms core-3-bdi1739.storm.ca [209.87.239.199]
 6 12 ms 19 ms 11 ms v502.core1.tor1.he.net [216.66.41.113]
 7 22 ms 22 ms 21 ms 100ge1-2.core1.nyc4.he.net [184.105.80.9]
 8 84 ms 84 ms 84 ms ny-paix-gni.twgate.net [198.32.118.41]
 9 82 ms 84 ms 82 ms 217-228-160-203.TWGATE-IP.twgate.net [203.160.22
 8.217]
10 82 ms 81 ms 82 ms 229-228-160-203.TWGATE-IP.twgate.net [203.160.22
 8.229]
11 82 ms 82 ms 82 ms 203.78.181.2
12 84 ms 83 ms 83 ms 203.78.186.70
13 84 ms * 85 ms 66.171.127.177
14 84 ms 84 ms 84 ms fortinet.com [66.171.121.34]
15 84 ms 84 ms 83 ms fortinet.com [66.171.121.34]

Trace complete.
```

In this scenario, the first hop contains the IP address `10.10.1.99`, which is the internal interface of the FortiGate. The second hop contains the IP address `172.20.120.2`, to which the `wan1` interface of the FortiGate is connected, so we can conclude that the route through `wan1` interface is being used for this traffic.

You can also see the route taken for each session by debugging the packet flow in the CLI. For more information, see [How to debug the packet flow](#).

## How to verify the correct firewall policy is being used

If you have more than one firewall policy, use the count column to check which policy is being used, the count must show traffic increasing. To do so, go to the **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6 Policy** page.

Also, debugging the packet flow in the CLI shows the policy ID that's allowing the traffic. For more information on debugging the packet flow, see [How to debug the packet flow](#).

## How to check the bridging information in transparent mode

When the FortiGate is set to transparent mode, it acts like a bridge and sends all incoming traffic out on the other interfaces. The bridge is between interfaces on the FortiGate.

Each bridge that's listed is a link between interfaces. Where traffic is flowing between interfaces, you can expect to find bridges listed. If you're having connectivity issues, and there are no bridges listed, that is a likely cause. Check for the MAC address of the interface or device in question.

## How to check the bridging information

To list the existing bridge instances on the FortiGate, use the following command:

```
diagnose netlink brctl list
```

### Sample output:

```
#diagnose netlink brctl list
list bridge information
1. root.b fdb: size=256 used=6 num=7 depth=2 simple=no
Total 1 bridges
```

## How to display forwarding domain information

You can use forwarding domains, or collision domains, in routing to limit where packets are forwarded on the network. Layer 2 broadcasts are limited to the same group. By default, all interfaces are in group 0. For example, if the FortiGate has 12 interfaces, only two may be in the same forwarding domain, which limits the packets that are broadcast to those two interfaces. This reduces traffic on the rest of the network.

Collision domains prevent the forwarding of ARP packets to all VLANs on an interface. Without collision domains, duplicate MAC addresses on VLANs may cause ARP packets to be duplicated. Duplicate ARP packets can cause some switches to reset. It's important to know what interfaces are part of which forwarding domains because this determines which interfaces can communicate with each other.

To manually configure forwarding domains in transparent mode, use the following CLI command:

```
config system interface
  edit <interface_name>
    set forward-domain <integer>
```

```
end
```

To display the information for forward domains, use the following command:

```
diagnose netlink brctl domain <name> <id>
```

where <name> is the name of the forwarding domain to display and <id> is the domain ID.

### Sample output

```
diagnose netlink brctl domain ione 101
show bridge root.b ione forward domain.
id=101 dev=trunk_1 6
```

To list the existing bridge MAC table, use the following command:

```
diagnose netlink brctl name host <name>
```

### Sample output

```
show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
Bridge root.b host table
```

port no	device	devname	mac addr	ttl	attributes
2	7	wan2	02:09:0f:78:69:00	0	Local Static
5	6	vlan_1	02:09:0f:78:69:01	0	Local Static
3	8	dmz	02:09:0f:78:69:01	0	Local Static
4	9	internal	02:09:0f:78:69:02	0	Local Static
3	8	dmz	00:80:c8:39:87:5a	194	
4	9	internal	02:09:0f:78:67:68	8	
1	3	wan1	00:09:0f:78:69:fe	0	Local Static

To list the existing bridge port list, use the following command:

```
diagnose netlink brctl name port <name>
```

### Sample output:

```
show bridge root.b data port.
trunk_1 peer_dev=0
internal peer_dev=0
dmz peer_dev=0
wan2 peer_dev=0
wan1 peer_dev=0
```

## How to check the number of sessions that UTM proxy uses

Each FortiGate model has a limit to the maximum number of sessions that the UTM proxy supports. The UTM proxy handles all the traffic for the following protocols: HTTP, SMTP, POP3, IMAP, FTP, and NNTP. If the proxy for a protocol fills up its session table, the FortiGate enters conserve mode, where it behaves differently, until entries and memory free up again.

### Conserve or failopen mode

Once you reach the limit, depending on the conserve mode configuration, no new sessions are created until old ones end. You can configure the behavior of the FortiGate when memory is running low or the proxy connection limit has been reached. There are two related commands for this in the CLI:

```
config system global
    set av-failopen-session {enable | disable}
    set av-failopen {idledrop | off | one-shot | pass}
end
```

To set the behavior for these conditions, you must enable `av-failopen-session`. When it's enabled, and a proxy for a protocol runs out of room in its session table, that protocol goes into failopen mode and behaves as defined in the `av-failopen` command.

`av-failopen` determines the behavior of the proxy until entries are free in the session table again for that proxy.

- **idledrop** — This option removes idle sessions from the session table, starting with the clients that have the most sessions currently open. This method assumes that idle sessions aren't being used and it won't cause problems to close these sessions. This is usually true, but some applications may be unable to open a session. If this occurs, try another method to check if this is really the problem. This is a secure option as no unscanned traffic is allowed to pass.
- **off** — This option turns off accepting any new AV sessions, but continues to process any existing AV sessions that are currently active. All the protocols listed (HTTP, SMTP, POP3, IMAP, FTP, and NNTP) are scanned by FortiGate Antivirus. If AV scanning is enabled, `av-failopen off` is selected, and the proxy session table fills up, which means that no new sessions of that type are accepted. For example, if the POP3 session table is full, and email AV scanning is enabled, no additional POP3 connections are allowed until the session table gets some free space. This is a secure option because no unscanned traffic is allowed to pass.
- **one-shot** — When memory is low, bypass the antivirus system. The term `one-shot` comes from the fact that once you're in `one-shot av-failopen` mode, you must set `av-failopen` to either `pass` or `off` to restart AV scanning. This is a very unsecure option because it allows all traffic without AV scanning, and it never reverts to normal without manual assistance.
- **pass** — When memory is low, bypass the antivirus system. The difference between `pass` and `one-shot` options is that when memory is freed up, the system automatically starts AV scanning. This is an unsecure option because it allows traffic to pass without AV scanning. However, it's better than `one-shot` because it automatically restarts AV scanning, when possible.

If the proxy session table is full for one or more protocols and your FortiGate enters into conserve or failopen mode, it appears as though the FortiGate has lost connection, network services are intermittent or don't exist, and yet other services work normally for a while until their sessions end and they join the queue of session-starved applications.



## How to examine the firewall session list

The firewall session list displays all sessions that the FortiGate has open. You can see if there are strange patterns, such as no sessions apart from the internal network, or all sessions are only to one IP address.

When you examine the firewall session list in the CLI, you can use filters to reduce the output. In the GUI, the filters belong to the interface.

### To examine the firewall session list - GUI

Go to **FortiView > All Sessions**.

### To examine the firewall session list - CLI

When you examine the firewall session list, there may be too many sessions to display. In this case, it's necessary to limit or filter the sessions displayed by source or destination address, or NAT'd address or port. If you want to filter by more than one of these, you need to enter a separate line for each value.

The following example shows filtering the session list based on a source address of 10.11.101.112:

```
FGT# diagnose system session filter src 10.11.101.112
FGT# diagnose system session list
```

The following example shows filtering the session list based on a destination address of 172.20.120.222:

```
FGT# diagnose system session filter dst 172.20.120.222
FGT# diagnose system session list
```

### To clear all sessions corresponding to a filter - CLI

```
FGT# diagnose system session filter dst 172.20.120.222
FGT# diagnose system session clear
```

## Check source NAT information

When you troubleshoot connections, remember NAT. NAT is especially important if you're troubleshooting from the remote end of the connection outside the firewall. On the **FortiView > All Sessions** list, pay attention to **NAT Source**, and **NAT Source Port**. These columns display the IP and port values after NAT has been applied.

Checking the NAT values can help you ensure that they are the values you expect, and that the remote end of the sessions can see the expected IP address and port number.

When you display the session list in the CLI, you can match the NAT'd source address (`nsrc`) and port (`nport`). This can be useful if multiple internal IP addresses are NAT'd to a common external-facing source IP address.

```
FGT# diagnose system session filter nsrc 172.20.120.122
FGT# diagnose system session filter nport 8888
FGT# diagnose system session list
```

## How to check wireless information

Wireless connections, stations, and interfaces have different issues than other physical interfaces.

### Troubleshooting station connection issue

To check whether a station entry is created on access control, use the following command:

```
FG600B3909600253 # diagnose wireless-controller wlaac -d sta
* vf=0 wtp=70 rId=2 wlan=open ip=0.0.0.0 mac=00:09:0f:db:c4:03 rssi=0 idle=148 bw=0 use=2
vf=0 wtp=70 rId=2 wlan=open ip=172.30.32.122 mac=00:25:9c:e0:47:88 rssi=-40 idle=0 bw=9
use=2
```

### Enable diagnostics for a particular station

This example uses the station MAC address to find where it's failing:

```
FG600B3909600253 # diagnose wireless-controller wlaac sta_filter 00:25:9c:e0:47:88 1
Set filter sta 00:25:9c:e0:47:88 level 1
FG600B3909600253 # 71419.245 <ih> IEEE 802.11 mgmt::disassoc <== 00:25:9c:e0:47:88 vap
open rId 1 wId 0 00:09:0f:db:c4:03
71419.246 <dc> STA del 00:25:9c:e0:47:88 vap open rId 1 wId 0
71419.246 <cc> STA_CFG_REQ(34) sta 00:25:9c:e0:47:88 del ==> ws (0-192.168.35.1:5246) rId
1 wId 0
71419.246 <cc> STA del 00:25:9c:e0:47:88 vap open ws (0-192.168.35.1:5246) rId 1 wId 0
00:09:0f:db:c4:03 sec open reason I2C_STA_DEL
71419.247 <cc> STA_CFG_RESP(34) 00:25:9c:e0:47:88 <== ws (0-192.168.35.1:5246) rc 0
(Success).
```

## How to verify connectivity to FortiGuard

You can verify connectivity to FortiGuard in the **Licenses** widget on the **Dashboard**. When your FortiGate is connected to FortiGuard, a green check mark appears for available FortiGuard services.

To verify connectivity to FortiGuard using the CLI, enter the commands `execute ping service.fortiguard.net` and `execute ping update.fortiguard.net`.

#### Sample output:

```
FG100D# execute ping service.fortiguard.net
PING guard.fortinet.net (208.91.112.196): 56 data bytes
64 bytes from 208.91.112.196: icmp_seq=0 ttl=51 time=61.0 ms
64 bytes from 208.91.112.196: icmp_seq=1 ttl=51 time=60.0 ms
64 bytes from 208.91.112.196: icmp_seq=2 ttl=51 time=59.6 ms
64 bytes from 208.91.112.196: icmp_seq=3 ttl=51 time=58.9 ms
64 bytes from 208.91.112.196: icmp_seq=4 ttl=51 time=59.2 ms

--- guard.fortinet.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 58.9/59.7/61.0 ms

FG100D# execute ping update.fortiguard.net
```

```
PING fds1.fortinet.com (208.91.112.68): 56 data bytes
64 bytes from 208.91.112.68: icmp_seq=0 ttl=53 time=62.0 ms
64 bytes from 208.91.112.68: icmp_seq=1 ttl=53 time=61.8 ms
64 bytes from 208.91.112.68: icmp_seq=2 ttl=53 time=61.3 ms
64 bytes from 208.91.112.68: icmp_seq=3 ttl=53 time=61.9 ms
64 bytes from 208.91.112.68: icmp_seq=4 ttl=53 time=61.8 ms
```

## How to perform a sniffer trace (CLI and packet capture)

When you troubleshoot networks and routing in particular, it helps to look inside the headers of packets to determine if they're traveling along the route that you expect them to take. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate has NP2, NP4, or NP6 interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP2, NP4, or NP6 interfaces, you should disable offloading on those interfaces.

## How do you sniff packets

Before you start sniffing packets on the CLI, you should prepare to capture the output to a file. A large amount of data may scroll by and you won't be able to see it without first saving it to a file. One method is to use a terminal program like puTTY to connect to the FortiGate CLI. Once the packet sniffing count is reached, you can end the session and analyze the output in the file.

The general form of the internal FortiOS packet sniffer command is:

```
diagnose sniffer packet <interface_name> <'filter'> <verbose> <count> <tsformat>
```

To stop the sniffer, type **CTRL+C**.

<b>&lt;interface_name&gt;</b>	The name of the interface to sniff, such as "port1" or "internal". This can also be "any" to sniff all interfaces.
<b>&lt;'filter'&gt;</b>	What to look for in the information the sniffer reads. "none" indicates no filtering, and all packets are displayed as the other arguments indicate.  The filter must be inside single quotes (').
<b>&lt;verbose&gt;</b>	The level of verbosity as one of:  <b>1</b> - print header of packets <b>2</b> - print header and data from IP of packets <b>3</b> - print header and data from Ethernet of packets <b>4</b> - print header of packets with interface name
<b>&lt;count&gt;</b>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run until you stop it with <b>&lt;CTRL+C&gt;</b> .

**<tsformat>**

The format of timestamp.

- a: absolute UTC time, yyyy-mm-dd hh:mm:ss.ms
- l: absolute LOCAL time, yyyy-mm-dd hh:mm:ss.ms
- otherwise: relative to the start of sniffing, ss.ms

For a simple sniffing example, enter the CLI command `diagnose sniffer packet port1 none 1 3`. This displays the next three packets on the port1 interface using no filtering, and verbose level 1. At this verbosity level, you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets and that 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diagnose sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]

0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface that are traveling between a computer with the host name of "PC1" and a computer with the host name of "PC2". With verbosity 4 and above, the sniffer trace displays the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type `CTRL+C`.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and icmp" 4
```

The following CLI command for a sniffer includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for example, PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

## How do you use packet capture

To use packet capture, the FortiGate must have a disk. You can enable the `capture-packet` in the firewall policy, using the following CLI commands:

```
config firewall policy
  edit <id>
    set capture-packet enable
  end
```

To configure packet capture filters, go to **Network > Packet Capture**.

When you add a packet capture filter, enter the following information and select **OK**.

<b>Interface</b>	Select the interface to sniff from the drop-down menu.  You must select one interface. You can't change the interface without deleting the filter and creating a new one, unlike the other fields.
<b>Max Packets to Save</b>	Enter the number of packets to capture before the filter stops.  This number can't be zero. You can halt the capturing before this number is reached.
<b>Enable Filters</b>	Select this option to specify filter fields
<b>Host(s)</b>	Enter the IP address of one or more hosts  Separate multiple hosts with commas. To enter a range, use a dash without spaces, for example 172.16.1.5-172.16.1.15, or enter a subnet.
<b>Port(s)</b>	Enter one or more ports to capture on the selected interface.  Separate multiple ports with commas. To enter a range, use a dash without spaces, for example 88-90
<b>VLAN(s)</b>	Enter one or more VLANs (if any).  Separate multiple VLANs with commas.
<b>Protocol</b>	Enter one or more protocols. Separate multiple protocols with commas. To enter a range, use a dash without spaces, for example 1-6, 17, 21-25.
<b>Include IPv6 Packets</b>	Select this option if you're troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.
<b>Include Non-IP Packets</b>	The protocols in the list are all IP based except for ICMP (ping). To capture non-IP based packets, select this feature. Examples of non-IP packets include IPsec, IGMP, ARP, and ICMP.

If you select a filter, you have the option to start and stop packet capture in the edit window, or download the captured packets. You can also see the filter status and the number of packets captured.

You can select the filter and start capturing packets. When the filter is running, the number of captured packets increases until it reaches the **Max Packet Count** or you stop it. When the filter is running, you can't download the output file.

When the packet capture is complete, you can download the \*.pcap file. You must use a third party application, such as Wireshark, to read \*.pcap files. This tool provides you with extensive analytics and the full contents of the packets that were captured.

To start, stop, or resume packet capture, use the symbols on the screen. These symbols are the same as those used for audio or video playback. Hover over the symbol to reveal explanatory text. Similarly, to download the \*.pcap file, use the download symbol on the screen.

For more information on troubleshooting with packet capture and packet sniffing, see ["Packet sniffing and packet capture" on page 25](#).

## How to debug the packet flow

Traffic should come in and leave the FortiGate. If you determine that network traffic isn't entering and leaving the FortiGate as expected, debug the packet flow.

You can only perform debugging using CLI commands. Debugging the packet flow requires that you enter a number of debug commands as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate has FortiASIC NP4 or NP6 interface pairs that are offloading traffic, this changes the packet flow. Before you perform the debug on any NP4 or NP6 interfaces, you should disable offloading on those interfaces. Enter `diagnose npu <interface pair> fastpath disable`, where `<interface pair>` can be `np4`, `np6`, `np4lite`, or `np6lite`.

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```
FGT# diagnose debug disable
FGT# diagnose debug flow filter add <PC1>
FGT# diagnose debug flow show console enable
FGT# diagnose debug flow show function-name enable
FGT# diagnose debug flow trace start 100
FGT# diagnose debug enable
```

The `start 100` argument in the above list of commands limits the output to 100 packets from the flow. This is useful to look at the flow without flooding your log or displaying too much information.

To stop all other debug activities, enter the command:

```
FGT# diagnose debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching security policy, and is in turn blocked by the FortiGate unit. The denied message indicates that the traffic was blocked.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825 msg="vd-root received a packet
(proto=6, 192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924 msg="allocate a new session-
013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597 msg="find a route: gw-
192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg=" Denied by forward policy
check"
```

# Troubleshooting resources

Fortinet provides customers with resources of valuable information about FortiOS technical issues, including:

## Technical documentation

Installation, Administration, and Quick Start Guides, as well as other technical documents, are available online at the [Fortinet Document Library](https://docs.fortinet.com): <https://docs.fortinet.com>

## Fortinet video library

The [Fortinet Video Library](https://video.fortinet.com) hosts a collection of video which provide valuable information about Fortinet products. <https://video.fortinet.com>

## Release notes

Issues that arise after the technical documentation is published will often be listed in the [Release Notes](#). To find these, go to the [Fortinet Document Library](#).

<https://docs.fortinet.com/fortigate/release-information>

## Knowledge base

The [Fortinet Knowledge Base](#) provides access to a variety of articles, white papers, and other documentation that provides technical insight into a range of Fortinet products. The Knowledge Base is available online at: <http://kb.fortinet.com>

## Fortinet technical discussion forums

An online technical forum allows administrators to contribute to discussions about issues that relate to their Fortinet products. Searching the forum can help an administrator identify if an issue has been experienced by another user. You can access the support forums at: <https://forum.fortinet.com/>

## Fortinet training services online campus

The [Fortinet Training Services Online Campus](https://www.fortinet.com/training.html) hosts a collection of tutorials and training materials which you can use to increase your knowledge of the Fortinet products. <https://www.fortinet.com/training.html>

## Fortinet customer support

You defined your problem, researched a solution, put together a plan to find the solution, and executed that plan. At this point, if the problem hasn't been solved, it's time to contact [Fortinet Customer Service & Support](#) for assistance. Prepare yourself by reading [How to work with Fortinet Support](#) on Fortinet's Cookbook site.

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>





**FORTINET®**



Copyright© (Undefined variable: FortinetVariables.CopyrightYear) Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.